Journal of Behavioral and Experimental Economics 108 (2024) 102151

Contents lists available at ScienceDirect

Journal of Behavioral and Experimental Economics

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

Who cares when Value (Mis)reporting may be found out? An Acquiring-a-Company experiment with value messages and information leaks

Daniela Di Cagno [a], Werner Güth [b], Tim Lohse [c], Francesca Marazzi [d,*], Lorenzo Spadoni [e]

[a] LUISS Guido Carli, Rome, Italy
[b] Max Planck Institute for Research on Collective Goods, Bonn, Germany
[c] Berlin School of Economics and Law, Berlin, Germany
[d] Tor Vergata University of Rome, Via Columbia 2, 00133 Rome, Italy
[e] Department of Economics and Law, University of Cassino and Southern Lazio, Cassino (FR), Italy

ARTICLE INFO

Dataset link: https://osf.io/4n8hy/, CC BY 4.0

JEL classification:
C78
C91
D83
D91

Keywords:
Acquiring-a-company experiments
Information leaks
Cheap talk
(Not) lying
Ultimatum bargaining

ABSTRACT

We modify the Acquiring-a-Company game to study how information leaks affect lying and market outcomes in an ultimatum bargaining setting with asymmetric information. Privately informed sellers send messages about the alleged value of their company to potential buyers. Via random leaks buyers, however, can learn the true value before proposing a price which the seller finally accepts or not. Only 14.5% of the messages are truthful, whereas two-thirds of all sellers exaggerate the company's value to persuade buyers to offer more, especially when the true value is small. Although a higher leak probability does not reduce the frequency of misreporting, it weakens overreporting and strengthens underreporting. Buyers who found out value misreporting anchor their price proposals on the true value but do not explicitly discriminate against liars. Sellers are mainly opportunistic and make their acceptances dependent on the resulting positive payoff. Even if ethical concerns do not seem to matter much, probabilistic leaks are welfare enhancing.

1. Introduction

One is usually reluctant to misreport private information when this can be revealed through information leaks. However, it is unclear whether the possibility of truth revelation completely eliminates misreporting or merely reduces its extent when the likelihood of information leaks increases. To sell at high prices, profit-maximizing sellers might exaggerate the value of what they offer for sale. But worse than selling too cheaply is not selling at all, e.g., due to being found out lying. We analyze the effects of lying in bargaining with privately value-informed sellers, appealing to many field situations and discussed in the literature of trade with private information since (Akerlof, 1978). Exploiting private information advantages can be questioned when so far uninformed buyers may learn the true values of sales items. In fact, this change in available information may alter the decision-making. Sellers would now have to consider that with overvalued offers they risk buyers not buying in the first place. Like obliging second-hand car dealers to reveal known but not easily recognizable damages of their cars, sellers might be held responsible for false statements or unrealistic promises if the buyers find out about their lies. In the field, where sellers may claim unawareness about the truth, intentional fraud may not be easily verifiable in court. Information leaks, when commonly expected, could instead alert potential trading partners and possibly avoid the need for legal regulation. One might even hope that the mere possibility of information leaks already prevents being exploited by privately informed sellers. If so, policy makers should encourage market transparency and revelation of information, e.g. by mandatory audits, in the course of which information may be disclosed, or even reward whistle blowing. This would boost efficiency-enhancing market exchanges even in the case of asymmetric information.

* Corresponding author.
E-mail address: francesca.marazzi@uniroma2.eu (F. Marazzi).

https://doi.org/10.1016/j.socec.2023.102151
Received 23 March 2023; Received in revised form 13 September 2023; Accepted 6 December 2023
Available online 7 December 2023

We capture such market situations via modifying the Acquiring-a-Company (AaC) experiment. In the AaC game, a seller owns a company that is evaluated higher by a potential buyer, with the two evaluations being perfectly (and linearly) correlated (Bazerman & Samuelson, 1983; Samuelson & Bazerman, 1985). Our experiment, therefore, relies on the stochastic ultimatum bargaining between an uninformed buyer who makes a price offer to a privately-informed seller, who can accept or not.[1] We make two modifications to the standard AaC game. First, before the buyer makes an offer, we let the privately informed seller send a numerical message regarding a company's alleged value that can be true or false. This allows for misreporting in two directions. Either the seller lies upwards and exaggerates the value of the company by an overreport to sell it for as much as possible. Or the seller lies downwards and understates it by an underreport reducing the potential payoff but increasing the chances of trade. Another modification we introduce is that, with a commonly known probability, the seller's private value information is leaked to the buyer. Hence, the buyer may not only become aware of the true value, but also whether and how the seller has been misreporting. Being so informed or not, the buyer proposes a price to acquire the company which the seller can accept (such that the company is sold) or reject.[2] Whereas at least commonly opportunistic trading partners would consider value messages as "cheap talk'', this could be questioned by monitoring via information leaks. If, for instance, the buyer would choose a very low price, when having found out that the seller has been lying by largely overstating the company value, this could prevent the seller from overstating value messages and, thus reduce overreporting. Additionally, moral concerns such as lying costs, guilt or shame may prevent sellers from lying. Stochastic leaks could also trigger social image concerns.

Trade is welfare increasing since sellers proportionally underevaluate the company. Across rounds, our experiment varies the commonly known proportional underevaluation parameter for the seller and, more importantly, also the probability of an information leak. Both, seller and buyer participants, are aware of both parameters when interacting across several successive rounds. Partners in the other role randomly change in order to discourage reputation effect. We refer to our modified Acquiring-a-Company experiment with information leak as *LAaC*.

Through our experiment, encompassing both stochastic and deterministic ultimatum bargaining, we aim to address the following research inquiries: What is the effect of a higher likelihood of information leaks on misreports by sellers? Does it increase the share of truthful reports or at least reduce the size of the lies? Or does it affect the direction of lying, i.e., over-versus underreporting? How do uninformed buyers react to value messages and how do informed buyers react to over- and underreporting? Do they reduce the offered price? What in turn determines sellers' acceptance of price proposals? Would they accept lower prices when aware of being found out lying? Finally, do information leaks enhance trade overall and thus welfare?

Our result are as follows. Most value messages are false (85.5%), but overreporting of company values accounts for only two-thirds of all cases. Instead, nearly one-third of these false value messages are underreports. Surprisingly, increasing the leak probability from 10% to 40% has only little effect on the frequency of truthful reports. What this variation in probability does alter, however, is the size of the lies and the structure of misreports. A higher leak probability reduces the average size of the lie (i.e. the difference between the true and the reported value), leads to a decrease in overreporting and a simultaneous increase in underreporting which can be considered evidence for sellers striving a positive social image as modest. Moreover, underreporting and truthtelling happen more often for higher values of the company, which increases the chances of trade. Overreporting is mostly concentrated at small to medium company values. Price proposals of buyers unaware of the true value anchor on the value message, while those who found out value misreporting anchor on the true values. However, they do not explicitly sanction liars but rather exhibit a certain inertia in suspicion: interacting with an overreporting seller in the past makes them propose lower price offers in future rounds. In contrast, sellers are fully opportunistic and make their acceptance decision mainly dependent on whether the resulting payoff is positive. Thus, morality concerns do not seem to matter much. Altogether, probabilistic leaks enhance trade and thereby also welfare.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 explains our modified Acquiring-a-Company game with information leaks. Section 4 presents the experimental protocol and states some behavioral predictions we want to test with our design. The experimental results are presented in Section 5. Section 6 concludes.[3]

2. Related literature

Morally it may already be questionable to not tell others that one privately knows something relevant, as studied by Dana et al. (2007) whose design employs a non-commonly known setup. Our research instead relies on a commonly known game form, namely a stochastic ultimatum bargaining setup with incomplete information. Strategic lying is mainly explored in modified ultimatum experiments by letting privately informed proposers lie about the pie size via pie-size messages (Mitzkewitz & Nagel, 1993). These could inform when to expect a more or less known share of truthful proposers and whether this in turn renders it worthwhile for opportunistic proposers to lie by overstating the pie size (Besancenot et al., 2013). Actually, pie-size messages become more deceptive the larger the pie (Vesely, 2014). Proposer messages can also concern the responder's outside option which should be respected via increased offers (Boles et al., 2000; Croson et al., 2003). Alternatively, responders may send messages about a (non-)favorable ECU-euro conversion rate[4] to induce higher offers by proposers (Koning et al., 2011). All these studies confirm substantial untruthfulness of messages but less deception when information is only indirectly transmitted via possibly information-revealing actions (Kriss et al., 2013).

In theory, pure cheap talk is ineffective (Kim, 1996) although it may matter behaviorally in experimental bargaining (Croson et al., 2003). However, when lies may be verifiable, messages are no longer cheap talk (Dato et al., 2019). Therefore, it is particularly interesting to investigate how probabilistic lying detection affects individual inclinations to send (un)truthful messages in bargaining and the interpersonal heterogeneity of such behaviors. So far the evidence indicates that proposers lie less in case of 50%-detection probability, compared to no detection at all (Anbarcı et al., 2015). Lowering the probability of detection weakens the effect. For example, a detection probability of even 25% does not prevent proposers from understating the pie size (Chavanne & Ferreira, 2017).

[1] The AaC-game model became first known as an elegant paradigm to analyze the winner's curse on bilateral monopoly markets with positively affiliated values of trading partners of which only one is perfectly aware. It has been largely overlooked that the AaC game actually seems to be the first stochastic ultimatum game in which only the responder is aware of the stochastic pie size when deciding (see Güth et al., 2023 for details). Another interesting feature of the AaC-model is its prediction of trade (when the surplus from trade is large and the seller accepts the buyer's price proposal) or no trade (when the surplus is too small) with the latter case resembling the prediction for lemon markets by Akerlof (1978).

[2] The leak probability alone allows to continuously connect the AaC-game and the ultimatum game.

[3] Appendix B contains the translated version of the experimental instructions.

[4] ECU standing for *e*xperimental *c*urrency *u*nit.

From a broader perspective, our analysis on the perceived (im)morality of lying when bargaining in asymmetric information settings also contributes to the literature on the so-called "dark side" of human nature, particularly lying on the one hand, and detecting lies on the other. The rapidly growing literature on lying behavior is mainly concerned with lies in which upward deviations from the truth occur in order to increase a player's payoff. In the well-known dice rolling experiments (Fischbacher & Föllmi-Heusi, 2013), for example, stating a number of 5 usually leads to the highest payoff. Players who roll a smaller number consequently have a monetary incentive to give a false higher indication. In contrast, lying behavior, in which lying is done in a downward deviation from the truth and therefore potentially reduces a player's payoff, still seems puzzling. Such downward lies, as Abeler et al. (2019) has coined them, are not yet part of recent theoretical models of lying costs (Abeler et al., 2019; Gneezy et al., 2018; Khalmetski & Sliwka, 2019) but first papers analyze possible motives for this. They argue that subjects may lie downward when a lower indication translates into a better reputation and signals higher honesty (Barron, 2019; Geraldes et al., 2021; Utikal & Fischbacher, 2013). Our setting provides a framework in which we can observe both lying for one's own maximum advantage (sellers overreporting the company value) and lies that result in payoffs lower than truthful statements (sellers underreporting the value). However, the motive for such downward lies is more complex in our case than in a standard die rolling task. While the payoff in the latter follows deterministically from the player's indication, in our LAaC experiment there is uncertainty about the seller's payoff. Thus, in addition to reputational concerns such as the desire to appear honest in order to maintain an honest self and social image, there is also an ambivalent monetary component because payoffs depend on whether trade occurs or not. On the one hand, high misrepresentations of the company's value would lead to a higher payoff in the event of trade. On the other hand, high reports of the company value can also make a potential buyer skeptical, as it is to be feared that they are exaggerated. In this case, a buyer may prefer not to buy at all. A truly high company value then poses a dilemma. Even an honest report (and even more so an exaggerated one) could appear as a lie to the buyer and consequently reduce the chances of trading. Against this background, a downward lie with regard to the company value can seem quite rational. For the seller, it lowers the payoff relative to a truthful report, but conversely, it increases the likelihood of trade. A detailed consideration of these aspects is new in the literature on lying.

Our work also contributes to the economic literature that deals with the detection of lies, which can basically occur in two ways. The first is accidental information leaks where, for example, a whistle blower makes a particular wrongdoing public. The threat of sanction is then thought to be an effective mechanism to prevent wrongful behavior in the first place. Whistle blowing has become recently investigated in experimental economics to the aim to assess its underlying intrinsic motivations (Reuben & Stephenson, 2013): how and through which channels do fines, leniency or rewards for reporting illicit activities affect cartel formation (Bigoni et al., 2012, 2015) and how effective are rewards for self-reporting bribery in the public sector (Abbink & Wu, 2017; Serra, 2012)? Differences in the social valuation of whistle blowers are equally important, as is the extent of harm caused by the misconduct (e.g. Butler et al., 2020). However, information that becomes public through such accidental leaks may not always be reliable and may not even be correct. Moreover, promoting whistle blowing can be a double-edged sword and, therefore, even backfire. The disclosure of private information on wrongdoings could lead to a shift in social norms so that, in the end, misconduct is no longer perceived as such (Bursztyn et al., 2020). The second type of information leak that leads to the discovery of lies is not random, but structured. This includes, for example, internal company monitoring procedures or externally initiated audits, in the course of which misconduct is uncovered and punished with a certain probability. The extensive literature on (tax) compliance analyzes such procedures. There is abundant experimental evidence on which factors impact misbehavior. As theorized in the seminal work by Allingham and Sandmo (1972), the level of the detection probability or the size of the penalty reduces misbehavior (see the meta-analysis of Alm & Malézieux, 2021). Face-to-face reports instead of anonymous ones via computer also lead to more honesty overall, although liars are then perceived as particularly honest and can thus hardly be detected (Dwenger & Lohse, 2019). Decision-making in groups rather than by individuals, in turn, leads to more dishonesty (Kocher et al., 2018; Lohse & Simon, 2021).[5] Our setting adds to this literature in terms of the positive effect of a reduction of asymmetric information. Reducing the advantages of lying may benefit a society by encouraging compliance and enhancing trade.

3. The Acquiring-a-Company game with information leaks (LAaC)

As in the basic Acquiring-a-Company setting, the only potential buyer B, when owning seller S's company, would evaluate it by $v \in (0,1]$. However, only seller S is aware of v whereas for buyer B the value v is uncertain and expected to be uniformly distributed in $(0,1]$, denoted by $v \in U(0,1]$. The seller, aware of B's expectations, evaluates the company via qv, where the parameter q with $0 < q < 1$ is commonly known. Seller S can send a value message $\hat{v} \in (0,1]$ to B which might be revealing, i.e., $\hat{v}(v) = v$ for all v, but also false, i.e. $\hat{v}(v) \neq v$ for $v \in (0,1]$. The total surplus from trade equals $(1-q)v$ and is always positive. Therefore, trade is always welfare enhancing.

The innovation of our setup is that private value information can be leaked and learned by B. With probability $w \in (0,1)$ buyer B, after receiving the value message $\hat{v}$ from the seller S but before proposing the price $p \in [0,1]$, may also learn the true valuation v, respectively qv. So with probability w the buyer would offer a price $p \in [0,1]$ to seller S, aware of both, the value message $\hat{v}$ and the true value v. This is denoted by $p = p(\hat{v}, v) \in [0,1]$. With the complementary probability $1-w$ buyer B only knows $\hat{v}$, the value message of seller S, when offering a price, denoted via $p = p(\hat{v}) \in [0,1]$. Finally seller S accepts or rejects the proposed price p.

For each $v \in (0,1]$ and given q with $0 < q < 1$, seller S earns $\delta(p)(p-qv)$ and B earns $\delta(p)(v-p)$. Note that both parties can lose; S obviously suffers a loss when accepting a price smaller than qv and B when a price proposal $p > v$, e.g., triggered by a value message $\hat{v} > p$, is accepted by S.

The benchmark solution for common(ly known) risk and loss neutrality and opportunism (in the sense of maximizing the own monetary payoff expectation) can be derived via backward induction. Optimal acceptance requires $\delta(p) = 1$ for $p \geq qv$ and $\delta(p) = 0$ otherwise. Anticipating this we first focus on the no-leak event. Here B expects to earn:

$$\int_{\frac{p}{q} \geq v > 0} (v-p)\,dv = \frac{p^2}{2q^2}(1-2q)$$

So the optimal price is $p^* = 0$ for $q > \frac{1}{2}$ and $p^* = q$ otherwise since $p^* = q$ is the lowest price guaranteeing trade for all $v \in (0,1]$. Our experiment considers both, $q > \frac{1}{2}$ implying no-trade prediction like Akerlof (1978) and $q < \frac{1}{2}$ which implies welfare enhancing trade with positive surplus $(1-q)v$ due to $v > 0$ and $q < 1$.

In case of the leak event with probability w, buyer B, aware of q, $\hat{v}$ and v, should offer prices $p^*(\hat{v}, v) = qv$ as a function of value v, i.e., $p^*(\hat{v}, v) = qv$ for all $v \in (0,1]$ where we abstract from the technicality that, in case of $p^*(\hat{v}, v) = qv$, seller S is indifferent between accepting and rejecting. So, with leak probability w the no trade result for all $v \in (0,1]$ and parameter $q > \frac{1}{2}$ is avoided. Like in non-embedded

[5] However, as (Fehr & List, 2004) show, there is also a risk that actual monitoring procedures may undermine trust and honesty and thus be counterproductive.

deterministic ultimatum games, there exist also other equilibria in weakly dominated response strategies but not in line with sequential rationality (Selten, 1975).

The LAac-game allows to assess behavior by both market sides. For sellers, we can study how $\hat{v}$ depends on v, how acceptance of price proposals p depends on $\hat{v}$ and v, and how both seller decisions are affected by the commonly known parameters q and w. It is also of interest to investigate if and how sellers' choices depend on the fact that a leak happened or not and whether, therefore, they have been found out lying.

Buyers, on the other side, confront two information conditions. First, one where the revealed level of v, together with its value message $\hat{v} = \hat{v}(v)$, let the buyer recognize the sign and size of possible misreporting.[6]

Second, the other where v is not revealed to the buyer and when seller and buyer should theoretically behave as in the AaC-setup, at least when both are known to be opportunistic.

If the buyer assumes the seller to be sufficiently lying averse and therefore believes in the truth of $\hat{v} = \hat{v}(v) = v$ for all $v \in (0,1]$, this would allow for value signaling and exploitative ultimatum pricing of buyers via $p = q\hat{v}$ for all $\hat{v} \in (0,1]$, i.e., the buyer would try to obtain the whole surplus from trade $\hat{v} - q\hat{v} = (1-q)\hat{v}$ for all $\hat{v} \in (0,1]$. There would be always trade, irrespective of $q \in (0,1)$, but the usual ultimatum exploitation. A buyer motivated by other drivers, instead, e.g., inequality concerns, might be willing to reward them by proposing (nearly) equal surplus sharing prices around

$$p(\hat{v}) = q\hat{v} + (1-q)/2\hat{v} = (1+q)/2\hat{v} \; for \; \hat{v} \in (0,1].$$

But the buyers expecting sellers to be sufficiently inequality averse could be naive. Denying any lying aversion would render value messaging as "cheap talk" and that the AaC-solution would not be questioned.

However, even when assuming that the buyer expects the seller to not at all mind lying, the additional modification of LAaC, namely that the buyer with positive leak probability w directly learns about the value v after receiving the value message $\hat{v}$ but before choosing p, could suffice to prevent lying and guarantee trade for all $v \in (0,1]$ via exploitative offers $p = qv$, irrespective of $q \in (0,1)$.

To illustrate this assume a positive leak probability w and that the seller, although not lying averse, is very afraid of being found out lying, for instance, due to fearing that this excludes all chances of further trade. Fearing and expecting that buyers refrain from trade when the seller is known to cheat would render value messages informative and guarantee trade for all $v \in (0,1]$ with the buyer demanding the whole surplus from trade. So both LAaC-modifications, messages and leaking v-information to buyers, could question the AaC-prediction (especially its no-trade one for $q > \frac{1}{2}$) and, thus, represent a first step in examining lying behavior by sellers in ultimatum bargaining.

In a nutshell, there are two reasons why sellers may not lie to the full extent when sending their value message. First, they may face moral concerns such as lying or guilt aversion, shame, or they may want to preserve the self-image of an honest person. Second, sellers may be afraid of ostracism, i.e. that once they have been found out lying, they might be sanctioned by no (or rather low) price offers by the buyers. In the literature all of these aspects have been studied before, but hardly ever by analyzing Bayesian Games whose commonly known priors determine how likely and in which ways such concerns matter for decision makers (in LAaC for sellers), respectively are expected to matter by trading partners (in LAaC by buyers).

[6] That actually features deterministic ultimatum subgames with commonly known pie size $(1-q)v$, the surplus from trade.

4. Design and predictions

4.1. Experimental protocol

Overview. The experiment is a parametrized version of the game explained in the previous section. Before the first round, half of all participants are randomly assigned to the role of a buyer while the others become sellers. Participants keep their roles throughout the 16 rounds of the experiment. At the beginning of each round, each buyer is randomly matched with one seller with whom to bargain about buying or not the seller's company.[7] Participants learned about their individual payoff at the end of each round. After the last round, everyone also had to answer a socio-economic questionnaire, which included a measure of self-reported risk propensity (see Dohmen et al., 2011).

Experimental task. A buyer and seller bargain whether to trade or not, i.e., whether the buyer acquires the company owned by the seller. The experiment lets the value v vary from 5 to 95 in increments of 5. First the computer randomly determines v and informs only the seller about it. However, the seller evaluates the company only by qv, where the seller's commonly known underevaluation parameter q is either 0.25 or 0.55.

The parameter q is crucial whether theoretically trade is predicted or excluded where the latter is related to Akerlof (1978). To illustrate this in more detail consider first of all the 0-probability of leak. This features the classic AaC game with cheap-talk value messages which theoretically do not matter. If both, seller and buyer, are known to be risk neutral, what we also induce experimentally, the theoretical prediction is "always trade", i.e., for all positive values at the price of q if q does not exceed $1/2$, respectively "no trade at all", i.e., zero prices, if q is larger than $1/2$. If instead a leak occurs, what results are ultimatum (sub) games with common value information in which the buyer acts as proposer and exploits the seller by offering the prices qv by which (s)he acquires – theoretically – all the surplus from trade. From this it follows that after the no-leak random event the resulting subgame is isomorphic to the classic AaC-game with common risk neutrality. So theoretically there is always trade with exploitative price offers after the leak random event and after the no-leak random event only when q does not exceed $1/2$ at the price q which is always acceptable for the seller, irrespective of the value. Behaviorally we neither expect exploitative price offers after the leak random event nor no trade at all after $q > 1/2$, possibly related to the winner's curse (see, for instance, Bazerman & Samuelson, 1983), and neither always trade for $q \leq 1/2$ and, if there is trade, always prices of q (see, for instance, Di Cagno et al., 2017).

In each round both, seller and buyer, know whether q is high or low, where both are equally probable. Then the seller, aware of v, sends a value message $\hat{v}$ to the buyer which may or may not reveal v to B. This value message is only subject to the same integer constraint as for v, i.e., the seller freely decides what to report. It is the only information of the buyer with probability $1-w$. However, and that is the innovation of the paper, with probability w (> 0), the buyer also learns the true v. The leak probability is either $w = 0.1$ or $w = 0.4$ and it is randomly selected across rounds. Knowing either only $\hat{v}$ or also v, the buyer then proposes a price p between 0 and 60 for the company which the seller can reject or accept. Payoffs are finally calculated and privately communicated to buyer and seller, respectively. In case of acceptance, the buyer earns the difference between the value of the company and the price, i.e., $v - p$, while the seller earns the difference between the price and the own evaluation of the company, i.e., $p - qv$. In case of a rejection, both earn zero. Fig. 1 summarizes the timing of events in one round, accompanied by an illustrative representation of these events through a game tree.

[7] We used random matching groups with six participants each, for a total of 17 such groups. Participants were not aware of such restricted matching.

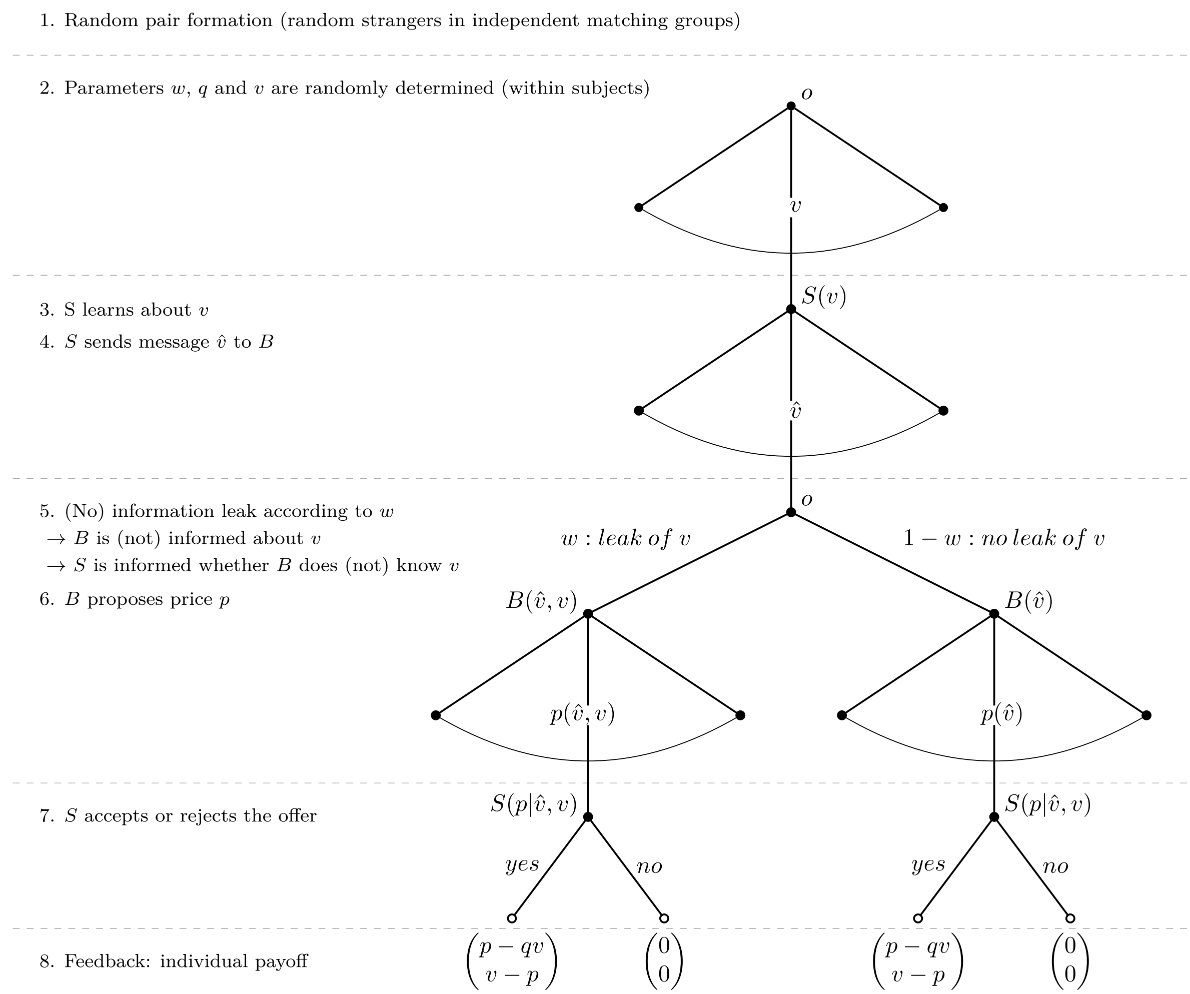

Fig. 1. Timing of events in one round.

Payoffs. After the experiment one round was randomly selected. A participant's earnings in this round were converted to probability points for gaining either EUR 4 or EUR 14 in a lottery. This lottery was played out on the participant's computer screen. Obviously such binary lottery incentives induce risk neutrality (see for an early study Roth & Malouf, 1979, and for an often misinterpreted criticism Selten et al., 1999).[8] Subjects earned an average of EUR 8.02 from the binary lottery plus a flat participation fee of EUR 6.

Implementation. Due to the constraints imposed by the COVID-19 pandemic, we conducted the experiment in a digital lab-like environment with payments administered immediately after the experiment via Prolific and Paypal.[9] A total of 102 subjects participated in 7 sessions with an average length of 105 min, including the time needed for connecting to the virtual cubicles, reading the instructions, the administration of the final questionnaire and feedback on earnings. The experiment was programmed in oTree (Chen et al., 2016) and carried out between July and October 2020 with the student participants of LUISS Cesare Lab recruited via ORSEE (Greiner, 2015). Students were from different fields of study, predominantly from economics, law and political science with an average age of 23.51 and 52% of them being female. No one participated in more than one session.[10]

Participants had one minute to take each of the three decisions throughout the experiment, i.e., value message, price proposal and acceptance, while between-decision feedback screens (e.g., feedback about leaks and earnings) lasted 30 s at most. In case a decision was not made in time, the computer selected a random decision in lieu of the participant and this was clearly stated in the instructions.[11] Moreover, a timer was shown in each of the decision screen. This protocol was adopted in order to avoid losing a full experimental session in case one of the participants would have dropped from the session (in which case a single matching group data would have been lost). The number of randomly-drawn choices was very small: 2 for the value message, 5 for the price proposal and 1 for acceptance, on a total of 1,632 choices (816 for each role). We removed these choices, and the corresponding interactions, from our data, which leaves us with 1,616 observations (808 for each role). Lastly, participants had the instructions available

[8] Specifically, even the lowest payoff does not rule out winning the larger reward, EUR 14, with positive probability and even the highest payoff does not guarantee this with 100% probability. So even when the buyer exploits the seller by offering $p(v) = qv$ after the w probability event, this would not deprive the seller of all chances to earn EUR 14. Similarly, even the largest possible loss will let the buyer earn the high reward EUR 14 with positive probability. What Selten et al. (1999) show is only that participants do not behave in line with expected utility theory even when rendering them risk neutral.

[9] For a description of the lab-like methodology see Buso et al. (2020).

[10] Summary statistics of the sample characteristics are reported in Table 1 in Appendix A.

[11] Moreover, if failing to communicate the value message or the price proposals, participants would have their probability of winning the high price reduced by 5% in case that specific round would have been selected for final payment.

during the whole experiment and, therefore, could refer to them at any moment; moreover, they could privately ask for clarifications from the experimenters at any time.

4.2. Behavioral considerations

From a behavioral perspective it is interesting to elicit lying and bargaining behaviors in one, overarching ultimatum experiment embedding both the stochastic and deterministic conditions, i.e. confronting participants with both privately and commonly known values v; this allows us to study how choices depend on the varying exogenously-induced leak probability.

In our setting, information leaks do not come with a monetary punishment for the seller who has misreported the value of the firm. However, the extensive literature on reputational concerns hints at the fact that being found out lying may run counter to one's desire to appear honest in the eyes of others and, thus, may harm one's social image (see, among others, Abeler et al., 2019; Bašić & Quercia, 2022; Fries et al., 2021 or Gneezy et al., 2018). In fact, the value of one's social image or reputation can serve as a powerful tool (Bénabou & Tirole, 2006): In general, it can have a direct impact, where individuals derive satisfaction from being seen in a favorable light, such as being perceived as honest or, more broadly, as a person of moral integrity. In our case, sellers may wish to project themselves as honest to ensure better price offers by buyers. In light of this, we expect sellers to reduce lying behavior when the probability of being found out, w, is higher. This can take two forms: either the frequency of truth messages increases, i.e. leaks completely discourage misreporting, or sellers still send untruthful messages but those are closer to the real message, i.e. the extent of the lie is smaller. We formalize our first behavioral hypothesis as follows

Leak Effects on Reporting (LEOR):

Information leaks may lead to (i) a higher share of truthful reports, (ii) a reduced size of the lies, i.e. $\hat{v}(v) - v$ *decreases.*

Since the seller's earnings are increasing in the price offered by the buyer, and this in turn depends on how valuable the buyer thinks the company is, the seller has an incentive to send a high value message. We consequently expect overreporting, especially when the true value of the firm is low. This is unsurprising. More interestingly, we also expect the opposite: underreporting. As already discussed in the introduction, a more recent strand of experimental literature finds that people might have a (reputational) incentive to lie downward. Although this often implies earning less in asymmetric information situations like our LAaC, downward lying is a way to gain a desirable image (Barron et al., 2022) and 'to appear honest' (Barron, 2019; Choshen-Hillel et al., 2020) in situations where the truth would seem like a lie because it is too good to be true.[12] In our experiment, sellers might resort to underreporting for two reasons: (i) to lower the chances that an uninformed buyer *perceives* them as dishonest, with possible repercussions on the price choice, and (ii) because of reputational concerns, since in case of an information leak the buyer *knows* whether the seller lied or not. We therefore expect sellers in our experiment to underreport the value of the firm, especially when the randomly assigned value of the firm is high, in order not to be perceived as liars. Also a high leak probability might have an effect on underreporting, as it comes with a higher chance to be found out lying and therefore sellers with reputational concerns have an incentive to make more buyer-favorable choices. However, this effect could also result in a higher frequency of truthtelling. The associated behavioral hypothesis is as follows

Underreporting (UR):

On the whole, participants tend to overreport. However, the share and extent of underreporting, i.e., $\hat{v}(v) - v < 0$*, can be substantial. We expect sellers to underreport more when* v *is high.*

The asymmetry in information and the impossibility of building a reputation for the seller, due to the stranger protocol, leave buyers with very little information to use when deciding about their price offer. The value message, $\hat{v}$, is the unique available information when a leak does not happen and it can serve as an anchor for buyers (see Croson et al., 2003; Tversky & Kahneman, 1974 or Meub & Proeger, 2015). Furthermore, the fact that a leak can happen might render cheap-talk messages more credible, due to the buyer anticipating the seller's reputational concerns. Therefore, fearing to state an offer which is inconsistent with the actual firm value, and thus to trigger a rejection from the seller, uninformed buyers might actually rely on the message received. On the other hand, when there is a leak and v is known, buyers operate in a deterministic setting and therefore will certainly anchor their price offer on the true value of the company. Our behavioral hypothesis for buyers is therefore

Leak Effects on Prices (LEOP):

When there is no information leak, uninformed buyers anchor on the message received, $\hat{v}$*, in setting the price. When buyers become informed, they take into account the true value of the firm,* v*.*

Anticipating buyers' reactions to information about the true value of the firm or to finding out that they were lied to is less straightforward. Once becoming informed about v, a buyer motivated by inequality concerns might propose a (nearly) equal surplus sharing whereas a rational one might simply demand almost the entire surplus and offer the seller peanuts. A similar effect can occur in case of negative reciprocity (Bolton & Ockenfels, 2000) when informed buyers react to having found out a lie via claiming a higher share of surplus. Furthermore, given the scarcity of information, buyers might rely on past interactions in spite of the randomly changing partner: good experiences with past sellers could increase the perceived trustworthiness, while bad ones might trigger suspicion.[13] The analysis of our experimental data will shed some light on these, possibly counteracting, effects.

Due to the leak probability, w, being positive but strictly lower than 1, the LAaC-setup encompasses both deterministic and stochastic ultimatum bargaining. Whereas the deterministic version of the ultimatum game is known for altruistic punishment – a major and influential finding where responders reject positive but still unfair offers (see Fehr & Gächter, 2002) –, the stochastic setting of the AaC-setup essentially crowds out altruistic punishing and rewarding of buyers by sellers (Bazerman & Samuelson, 1983; Samuelson & Bazerman, 1985 and relatedly Angelovski et al., 2020 and Güth et al., 2023): sellers nearly always accept price offers which result in some positive payoff for them. We therefore expect sellers to mostly behave rationally and thus to accept price offers that yield a gain and reject otherwise. The reason seems to be that in crucially stochastic environments random interpersonal payoff comparisons are rather cumbersome.

The associated behavioral hypothesis is

Seller Opportunism (SO):

We expect sellers to behave opportunistically, i.e. to accept when $p > qv$ *and to reject when* $p < qv$*.*

To sum up, if our hypotheses find confirmation in the analysis of observed behavior, we should find that information leaks discourage lying on the seller's side. Buyers, on the other hand, could offer a price which is more coherent with the value of the firm, which is either known or proxied by a cheap-talk message, with the latter rendered more credible by the possibility of the leak. Therefore, we should observe that, when the probability of a leak is higher, trade is more

[12] Although in a different context, Dana et al. (2007) show experimentally that one often accepts minor losses when these allow to hide own misbehaving.

[13] Buyers gain experience on sellers' reporting behavior in two possible ways: either when a leak happens or in case their offer is accepted. In this last case, in fact, they can infer whether the value message was truthful or not when they are informed about their payoff for that round.

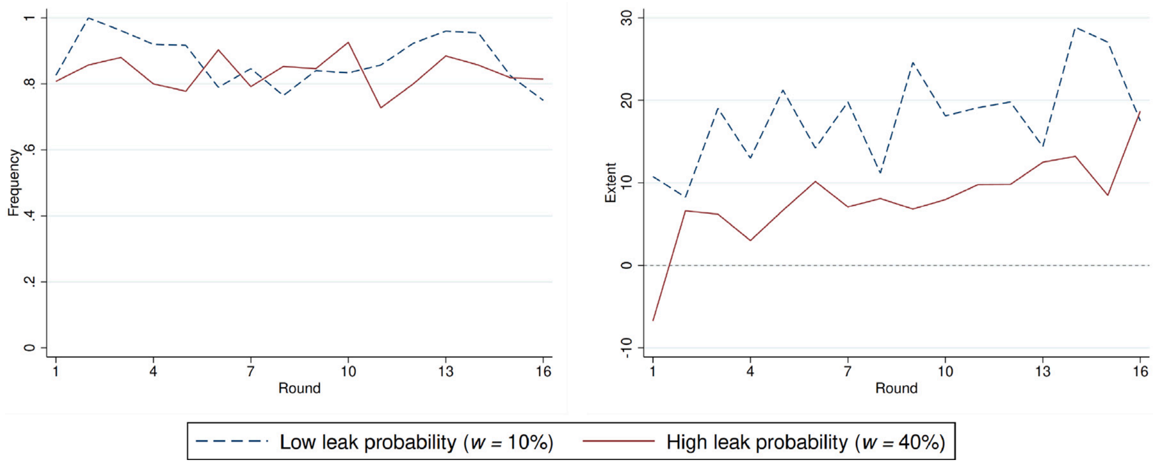

Fig. 2. Dynamics of misreporting behavior: misreporting frequency (left panel) and size of lie (right panel).

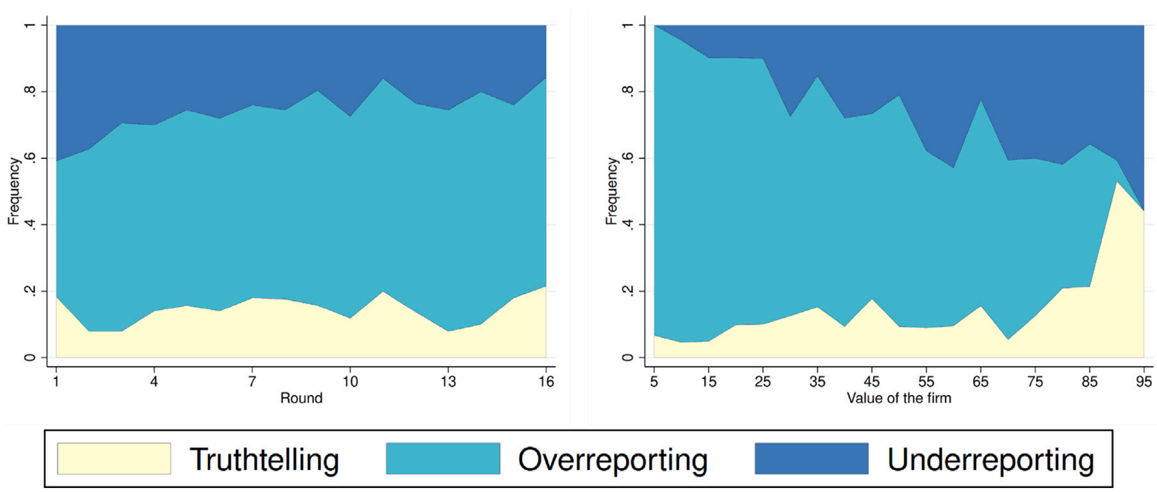

Fig. 3. Truthtelling, over- and underreporting frequencies by round (left panel) and by company value (right panel).

frequent. Reminding the reader that in the AaC game trade is always welfare enhancing, we state our last hypothesis:

Welfare Increase through Leaks (WITL):

Information leaks increase the likelihood of trade which in turn increases welfare.

5. Experimental results

In this section we analyze the sellers' (mis)reporting behavior and the variables that determine the market outcome such as buyers' price proposals and sellers' decision to accept or reject the proposals. In each case, we first present descriptives before performing regression analyses.

5.1. (Mis)reporting behavior

5.1.1. Descriptives

Overview & Dynamics. Truthful value messages are rare (14.5%) since most sellers lie and misreport their company's value (85.5%). The average size of a lie (measured as $\hat{v} - v$) is 12.67 (standard deviation: 27.13).[14]

Fig. 2 shows the dynamics. The frequency of misreporting is rather time invariant on a comparably high level and does not depend on the leak probability (left panel).[15] This is somewhat at odds with the *LEOR(i)*. However, the size of sellers' lies increases substantially across rounds (right panel). More interestingly, there is a stable gap in the size of the lies (apart from the very last round): the average difference between the reported and the true value of the company is smaller when the leak probability is high. This is a first piece of evidence for the *LEOR(ii)*. In sum, more probable leaks do not lead to fewer dishonest reports, but to less severe ones.

Around 70% (69.9% to be precise) of all untruthful reports result from sellers overreporting v, with a maximum lie of 90. Interestingly, in the case of the remaining almost third (30.1%), the reported value, $\hat{v}$, is lower than the real v with a minimum of −65. Thus, a large number of sellers actually underreport. Fig. 3 (left panel) illustrates the frequency and dynamics of under- and overreporting (along with truthtelling) across rounds. Despite the clear and increasing dominance of overreporting, both underreporting and truthtelling still occur to a significant extent in later rounds. Underreporting decreases after the first three rounds, meaning that some participants are learning in which direction it is more convenient to lie.

[14] We did not impose additional social norms of truthtelling in the instructions which could have resulted in lower overall lying rates. However, the effect of the parameter constellation that we are examining here would be independent of this.

[15] A paired t test run on independent matching-group level frequencies of misreporting by w confirms a non significant difference (p-value = 0.164).

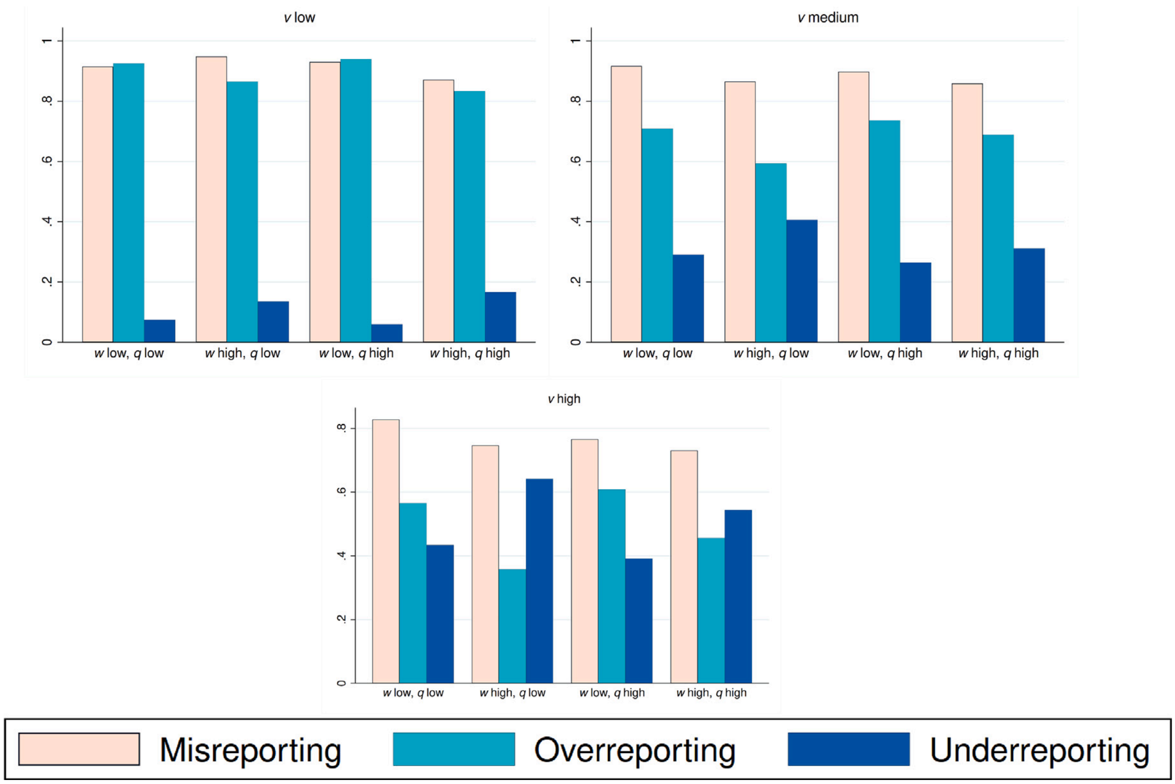

Fig. 4. Frequency of misreporting and percentage of over and underreporting by parameter constellations.

Value of the company. To investigate the reasons behind these different types of misreporting, Fig. 3 (right panel) illustrates their dependence on the company value v rather than time. Underreporting and truthtelling increase with value v. Overreporting is mostly concentrated at small to medium size values v. We summarize our findings which provide evidence for *UR* in

Result 1:

Misreporting in value messages is massive (about 86%) and dominated by nearly 60% overreporting, especially in the lower value range of v. However, under- and truthfully reporting are substantial, too.

Parameter constellations. Next we investigate if and how (mis) reporting behavior changes with parameter constellations. In order to take into account the role of the value of the company, we split the distribution of v in three categories: *low* if $v \in \{5, 10, 15, 20, 25, 30\}$, *medium* if $v \in \{35, 40, 45, 50, 55, 60\}$ and *high* if $v \in \{65, 70, 75, 80, 85, 90, 95\}$.

Fig. 4 attempts to visualize the joint effects of the parameter constellation for the leak probability, w, and the extent of the undervaluation of the company, q, as well as of the value of the company, v. With respect to the latter, the data corroborates our previous findings: the frequency of underreporting increases as v increases whereas the frequency of overreporting decreases. Surprisingly, the proportion of untruthful versus truthful reports remains rather stable for every combination of parameters, and an increase in w by factor 4 only leads to a small increase of truthful reports. What does change drastically, however, is the structure of the misreports. Subjects seem to substitute overreporting with underreporting, instead of truthtelling. To be more precise, a higher leak probability always leads to a decrease in overreporting and a simultaneous increase in underreporting (compare the first to the second pillars, and the third to the fourth pillars, respectively). Such behavior can be seen as evidence that sellers strive for a positive social image as modest in sharing the pie which only materializes with a high leak probability. They seem to want to achieve this by understating the value of the company and thus enabling the buyer to make a particularly favorable acquisition. The fact that sellers attach a monetary value to their social image is revealed by the fact that they are willing to accept a lower price and, thus, payoff, than would have resulted from a value message with the true value of the company. The increase in underreporting for a higher leak probability is stronger, the higher the value v. This reveals a second motive for downward lying: the already known danger of scaring off the buyer with honest (and even more so exaggerated) value messages and thus not achieving a trade.

In contrast, the effects for the extent of the underevaluation of the company, q, are less pronounced. Only for medium and high values v and a high leak probability, we observe a mild increase in overreporting along a corresponding decrease in underreporting (compare the first to the third pillars, and the second to the fourth pillars, respectively).

5.1.2. Regression analysis

Our findings from the descriptive data are supported by regression analysis. We apply a three-level, random intercept model with observations nested at the individual and matching group levels. Table 1 shows the results where the dependent variable is the seller's (mis)reporting extent, i.e. $\hat{v} - v$.

Our set of explanatory variables includes the percentage extent of misreporting in the previous round (measured as the difference between the reported and the real company value relative to the real one), a dummy variable regarding the detection of a lie in the previous round and a dummy variable regarding the acceptance of the proposed price and thus the conclusion of the trade. Furthermore, we use the three-way division of the company value explained above (a medium v is the reference category), and dummy variables for the parameters q and w, which are each 1 if the parameter values are high (i.e. 0.55 and 0.4, respectively). We also control for the subjects' risk propensity. The set of controls is complemented by round dummies and demographic controls, which include the gender of the subjects, their age, whether they are studying economics, whether they are experienced in participating in experiments and from which Italian macro-area they come (i.e. dummy for center of Italy, where the university is located). Table

Table 1
Determinants of sellers' (mis) reporting.

Depvar: extent of misreporting, $\hat{v} - v$			
	(1)	(2)	(3)
$\frac{\hat{v}-v}{v}$ at $t-1$	0.579**	0.613**	0.583**
	(0.250)	(0.251)	(0.249)
Lie detected ($t-1$)	−0.678	−1.546	−0.875
	(1.460)	(1.536)	(1.458)
Accepted ($t-1$)	3.674**	3.907***	3.673**
	(1.485)	(1.483)	(1.480)
Baseline: medium v			
Low v	19.919***	19.899***	19.886***
	(1.506)	(1.501)	(1.501)
High v	−17.554***	−17.551***	−17.399***
	(1.489)	(1.486)	(1.486)
q is high	4.074***	3.938***	4.032***
	(1.209)	(1.207)	(1.205)
w is high	−9.855***		−18.054***
	(1.232)		(4.183)
Baseline: no change in w (low)			
w was high, now low		4.333**	
		(1.851)	
w was low, now high		−7.496***	
		(1.788)	
No change (high)		−7.225***	
		(1.895)	
Risk propensity	3.319***	3.263***	2.652**
	(0.998)	(0.999)	(1.056)
Interaction w high ∗risk prop.			1.213*
			(0.645)
Demographics	✓	✓	✓
Round dummies	✓	✓	✓
Observations	759	759	759
Number of (matching) groups	17	17	17

Notes: standard errors in parentheses, $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Estimates from three-level hierarchical model with random effects, with observations nested at the individual and matching group levels. All specifications include round dummies and demographic controls (gender, age, geographic origins (with students coming from the central regions of Italy as a baseline), a dummy for students in economics and the level of experience in lab experiments).

1 in Appendix A reports the summary statistics of such demographic controls.

Overall, misreporting is path dependent as the extent of a lie in the previous round has a robust (across specifications) explanatory power. In contrast, whether or not the seller was caught lying in the previous round due to the leak event has no effect on the seller's current misreporting behavior. Interestingly, a seller who had accepted the proposed price previously, is likely to misreport higher in the current round. As expected, the value size v plays a predominant role: sellers overreport when v is small, and underreport when v is large. However, this is, at least in part, structural because the range of possible overreporting $\hat{v}$ is small if v is already high, and likewise it is also small for underreporting $\hat{v}$ if v is small. Importantly, and as already revealed by the descriptive statistics, the high leak probability decreases the extent of misreporting. This supports the idea of the *LEOR(ii)*. In turn, a higher extent of underevaluation, q, increases the extent of misreporting, as sellers anticipate they will have lower profits in this case.

To further explore the role of w, model (2) investigates the effect of a 'change in scenario', i.e. whether the leak probability which is randomly determined in every round has increased or decreased (or remained the same), compared to the previous round. Compared to a low w in both rounds, high w triggers lower misreporting extent, independently of its value before, while a decrease in w encourages it.

Despite experimentally induced common(ly known) risk neutrality, the post-experimentally self-assessed risk propensity (the higher the value, the more risk loving the subject is) could play a role.[16] More risk loving participants misreport more and the interaction of risk propensity with leak probability in model (3) is weakly significant.[17] This, of course, shows that uncertainty inclinations reflect much more general disposition than narrowly captured by cardinal utility theory and the expected utility hypothesis of expected utility maximization.[18]

Finally, to complement the analysis on misreporting, Table 2 in Appendix A reports the results of a random-effects multinomial logit regression where reporting behavior is categorized as overreporting, truthtelling and underreporting. Besides confirming what we already observed in Table 1, the estimates corroborate the visual impression of Fig. 4. A higher leak probability increases truthtelling, but especially underreporting while, in turn, overreporting decreases. Underreporting and truthtelling are also more frequent for higher company values. Taken together, this may indicate that sellers seem to strive for a positive social image as modest and are eager to increase chances of trade.

We summarize our findings which provide evidence for *LEOR(ii)* and *UR* in

Result 2:

An increase in the leak probability reduces the size of the lies and changes the structure of misreports with underreporting partially crowding out overreporting. Furthermore, the extent of misreporting is partly path dependent due to inertia in misreporting. In turn, a higher extent of underevaluation, q, increases the extent of misreporting.

5.2. Market outcomes

5.2.1. Descriptives

Table 2 reports summary statistics for buyers' behavior in the bargaining process, overall and differentiated according to whether v was private knowledge (upper part *a*) or public knowledge due to an information leak (lower part *b*). In the latter case, we additionally distinguish whether a lie has been detected. Together with average price proposals p, we show the implied share of the value (hereafter shv) and of the surplus (hereafter shs, with $s = v - qv$) stemming from such proposals. In particular, shv is defined as $\frac{v-p}{v}$ when there is an information leak and v is known; when there is no leak, its equivalent is $sh\hat{v} = \frac{\hat{v}-p}{\hat{v}}$. Similarly, the proposed sharing of the surplus, shs, is defined as $\frac{v-p}{(1-q)v}$ when there is an information leak and v is known and $\frac{\hat{v}-p}{(1-q)\hat{v}}$ when there is no leak.[19]

Proposed prices are slightly higher when v is not leaked (*a*) compared to when there is an information leak (*b*), and the difference is only weakly statistically significant according to two-sample t test run on matching-group average observations, with p-value = 0.0504. What is consistently different, however are the requested shares of value. In the first case buyers aim at earning 47% of the (hypothetical, as they rely on $\hat{v}$) pie size while in the latter when they are informed they lower their requested shv to only 30%. With such low requests, buyers seem to want to reduce the probability that the seller might reject their price proposal. In addition, the higher share of $\hat{v}$ demanded when v is not

[16] It is important to note that the commonly used question for risk assessment, 'Are you generally a person who is fully prepared to take risks or do you try to avoid taking risks?' does not mention probabilities or monetary rewards, therefore it is more general than the neoclassical definition of risk aversion.

[17] The demographic controls do not exhibit significant effects in predicting lying. Coefficients for demographics and round dummies are available upon request.

[18] Note that this questions expected utility theory and not that, according to expected utility theory, binary-lottery incentives induce risk neutrality.

[19] As the definitions of shv and shs differ depending on whether there was a leak or not, we do not report their averages for the full sample.

Table 2
Summary statistics of buyers' price and value/surplus sharing proposals.

	Buyers' proposals					
	p	Share of value		Share of surplus		N
		$\hat{v}$	v	$\hat{s}$	s	
All choices	29.71	–	–	–	–	808
(a) v not leaked	30.23	0.47		0.83		600
(a.1) overreporting	32.18	0.51		0.90		355
(a.2) truthtelling	33.33	0.44		0.79		87
(a.3) underreporting	24.15	0.41		0.71		158
(b) v leaked	28.20		0.30		0.55	208
(b.1) v leaked, lie detected	27.49		0.27		0.51	177
(b.2) v leaked, lie not detected	32.26		0.44		0.80	31

Notes: share of value and of *surplus* are defined as the payoff the buyer proposes for herself over the available amount (i.e. the value or the surplus), conditioned on the information she has. When v is not leaked, $sh\hat{v} = \frac{\hat{v}-p}{\hat{v}}$ and $sh\hat{s} = \frac{\hat{v}-p}{(1-q)\hat{v}}$; when v is leaked, $shv = \frac{v-p}{v}$ and $shs = \frac{v-p}{(1-q)v}$.

known might signal that buyers are anticipating sellers' overreporting tendency.[20]

A distinction based on the type of the lie when v is not leaked (cases *a.1* to *a.3*) and whether or not it was detected in the leak event (cases *b.1* and *b.2*) refines the picture. When v is not leaked, sellers' overreporting and underreporting have, of course, opposite effects on price proposals with the former being significantly higher than the latter. Nonetheless, the requests on shv and shs vary little in the type of the lie. Lastly, in case of leaks, i.e. when there is ultimatum bargaining with complete information like in usual ultimatum games, price proposals are lower when a lie was detected (case *b.1*; 27.49) compared to when the seller sent a truthful message (case *b.2*; 32.26). This statistically significant discount (paired t test on matching group-level averages; p-value=0.0402) could be considered as evidence for buyers discriminating against lying sellers. However, given that sellers are more likely to overreport, overall and more specifically for low v-values, it seems much more reasonable to conclude that in this last case buyers simply adapt their pricing behavior to the known v. Put differently, there is no evidence that buyers discriminate or punish liars. They rather adapt price proposals to their information set. When only $\hat{v}$ is known, they anchor to this value message and prices are higher because sellers are more likely to overreport. As a matter of fact, when a lie is detected buyers claim a lower share of v (case *b.1* in Table 2) compared to when they are not informed (case *a*). This hints to the fact that they claim more when uninformed because they are anticipating that $\hat{v}$ is overreported. However, the low number of observations does not allow for further robust investigations.

Table 3 presents summary statistics for sellers' propensity to accept. In Panel A we report the number of trade opportunities (N), the number of acceptance decisions (α, i.e. how many of the opportunities turned into realized trades) and the acceptance rate (α/N) overall, by information condition depending on whether v was leaked or not, and by q-value. In Panel B we report the same quantities distinguishing whether acceptance would imply a positive, null or negative payoff for the seller.

Overall, sellers accept 75.1% of the price proposals; however, if v is leaked (case *a.2*), then acceptance occurs significantly more frequently than if v is not leaked (case *a.1*): 80.77% versus 73.17% (paired t test run on acceptance rate at the matching group level, p-value = 0.017). Furthermore, the acceptance rate is significantly higher when q is low (case *b.1* versus case *b.2*; paired t test on matching group-level averages; p-value = 0.000). However, this follows somewhat straight from the seller's payoff structure. The wider the gap in the valuation of the company between the buyer and seller, the wider the range of price proposals that the seller can accept.

Table 3
Summary statistics of sellers' acceptance, overall (Panel A) and by potential payoff sign (Panel B).

Panel A			
	N	α	*% accepted* (α/N)
All choices	808	607	75.12%
(a.1) v not leaked	600	439	73.17%
(a.2) v leaked	208	168	80.77%
(b.1) q is low	409	343	83.86%
(b.2) q is high	407	270	66.34%

Panel B				
	N	α	*% accepted* (α/N)	*% of sample* (α/808)
Positive payoffs	605	554	91.57%	68.56%
Null payoffs	10	6	60.00%	0.74%
Negative payoffs	193	47	24.35%	5.82%

Notes: N indicates the absolute frequency of cases irrespective of the seller's decision (i.e. the number of potential trade opportunities) while α indicates in how many of those cases the seller accepted (i.e. the number of the realized trades.

When distinguishing choices according to the sign of the potential payoff (Panel B), sellers nearly always accept buyers' price proposals if this yields a positive payoff and reject if negative. This is clearly in line with sellers being opportunistic (*SO*). Quite surprisingly, in nearly 25% of the cases when acceptance would cause a loss (corresponding to 5.82% of all observations), sellers accept nevertheless and get a negative profits. Compared to this, sellers rejecting positive profits is more rare (around 8.43%). This suggests some altruistic rewarding and sanctioning, respectively.[21]

Finally, Fig. 5 illustrates the price dynamics. When buyers are uninformed (v is not leaked; left panel, dashed line), average prices are more volatile than when v is leaked (solid line), although this may also be due to fewer observations without v-leaks. Overall, we see a stable decline of price proposals during the last 6 rounds in both cases. This could be due to late learning of the winner's curse, i.e., of realizing that for large q, buyers' expected payoffs from trade are negative. Only for low q their losses would be overcompensated by gains. Despite the declining price offers, acceptance rates are relatively stable at around 75% over the course of the rounds (right panel). Of course, acceptance rates are higher for low q.

[20] A similar picture emerges for the share of the surplus, shs, which is either 83% or 55%. However, the fact that the requested shs amounts to 83% when v is not known suggests that participants fail to anticipate the effect of q on the aggregate payoff.

[21] This would somehow contradict the effect *SO*. However, the corresponding number of acceptance choices are 47 and 51, respectively, which is too low for reliable behavioral conclusions.

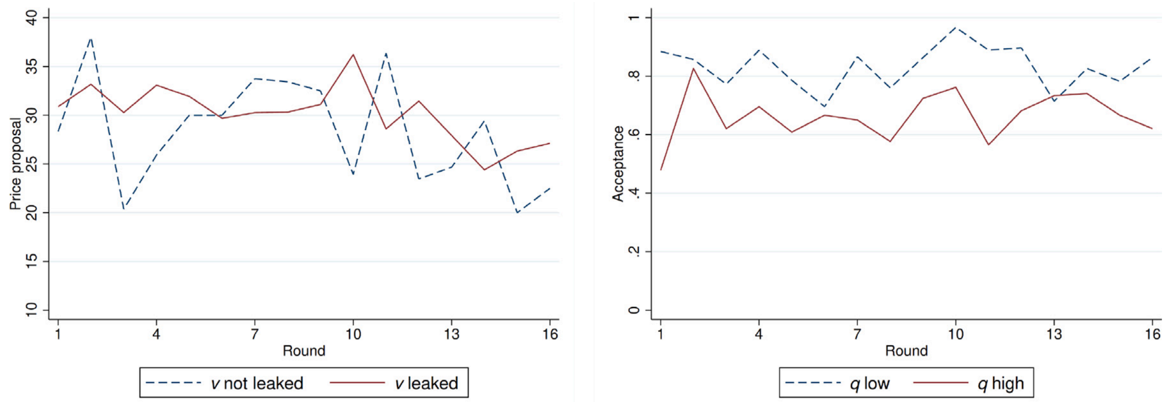

Fig. 5. Dynamics of buyers' price proposals (left panel) and sellers' acceptance (right panel).

5.2.2. Regression analysis

Table 4 presents the results of regression analyses with buyers' price proposals, p, as dependent variable. We run similar, but not identical, analyses for the split sample based on whether there has been an information leak or not (first two columns) as well as for the full sample (third column). Given that the sample split originates two unbalanced panels with a smaller number of observations than the full one, we resort to a dynamic random effects approach. To account for matching-group specific effects, all regressions include matching group dummies, as well as round dummies and the set of demographics described in Section 5.1.2.

When the buyer is uninformed since v was not leaked, the strong significance of the previous price proposal, p_{t-1}, reveals a clear path dependence in pricing behavior. Interestingly, buyers' price proposals also depend on their experience with sellers' misreporting behavior in the previous round. We categorize misreports as the relative difference between a past value message and the real value. Underreporting results in a negative difference. In contrast, modest overreporting is reflected in a positive difference below 100% whereas excessive overreporting comes along with a positive difference above 100%. The reference category is truthful reporting.[22] The analysis reveals that buyers substantially decrease their price offers after having met a seller who overreported compared to having met truthful sellers. So there is inertia in suspicion.

In addition, estimates show that not knowing v let buyers base their decision mostly on the value message received, $\hat{v}$. Obviously, in case of an info leak, it is not $\hat{v}$ but rather v itself which has explanatory power for the price proposals (third column in Table 4). We therefore find confirmation of *LEOP*. Lastly, the dummy variable *detected*, which is equal to 1 in case v is known and the seller misreported it, turns out to be not significant. Thus, we find no evidence for buyers discriminating against lying sellers. When there is an information leak, buyers rather seem to condition their pricing behavior much more to the real value v.[23]

Result 3:

When there is no information leak, price offers depend on the message received and own past choice. When there is an information leak, buyers anchor on v but they do not discriminate against lying sellers. Furthermore, there is inertia in suspicion, in the sense that interacting with an overreporting seller decreases later price offers.

[22] Since the parameters q and w are highly correlated with lying behavior, we avoid including them in the regression given the presence of the value message and of the lie detection variable.

[23] Across all three specifications, among the demographic controls, females offer significantly higher prices whereas risk propensity turns out insignificant. Coefficients for demographics and matching group and round dummies are available upon request.

Table 4
Determinants of buyers' proposed price.

	Depvar: price offer, p_t		
	v not leaked	v leaked	Overall
p_{t-1}	0.241***	0.063	0.218***
	(0.041)	(0.052)	(0.038)
Reporting at $t-1$, baseline: truthtelling			
Underreporting	1.236	−2.307	0.953
	(1.349)	(2.174)	(1.465)
Modest overreporting	−1.960	−4.061**	−1.488
	(1.438)	(2.018)	(1.269)
Excessive overreporting	−3.349**	−4.924**	−3.481**
	(1.365)	(2.403)	(1.417)
$\hat{v}$	0.341***		0.323***
	(0.030)		(0.029)
v		0.480***	
		(0.041)	
Detected		−1.469	−1.608
		(2.335)	(1.465)
Demographics	✓	✓	✓
Round dummies	✓	✓	✓
Matching group dummies	✓	✓	✓
Observations	554	197	751

Notes: Robust standard errors in parentheses, $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Estimates from random effects model. All specifications include matching group and round dummies and demographic controls (gender, age, geographic origins (with students coming from the central regions of Italy as a baseline), a dummy for students in Economics and the level of experience in lab experiments). Categories of reporting at $t-1$ correspond to $\frac{\hat{v}-v}{v}$ being either negative (underreporting), null (truthtelling, baseline category), <100% (modest overreporting) or ≥100% (excessive overreporting).

Finally, Table 5 presents the marginal effects of pooled probit regressions with sellers' acceptance as dependent variable. We again show results separately for the cases in which v is (not) leaked (first two columns) and for the full sample (third column). All regressions include round and matching group dummies and the same set of demographic controls. In Panel A we use as regressor the price offer received; in Panel B we include the seller's lie (as percentage of v) to check whether own lying behavior affects acceptance.

As predicted, sellers are fully opportunistic (in line with the idea of *SO*). The main determinant of acceptance is whether the resulting payoff is at least positive, i.e. whether the proposed price p is larger or at least equal to the seller's valuation of the company, qv. With the exception of an info leak, the price level itself also has a significant, albeit small, effect. Massively overreporting also enhances acceptance, compared to truthtelling, and actually triggers higher and thereby more acceptable price offers. But this phenomenon disappears when restricting the analysis to cases when informed buyers offer prices. Moreover, sellers do not discriminate against buyers who they have

Table 5
Determinants of sellers' acceptance.

Depvar: acceptance decision			
Panel A			
	v not leaked	v leaked	Overall
Non-negative payoff	0.384***	0.370***	0.367***
	(0.019)	(0.039)	(0.020)
p_t	0.004***	0.002	0.004***
	(0.001)	(0.002)	(0.001)
Detected		−0.012	0.024
		(0.085)	(0.027)
Panel B			
	v not leaked	v leaked	Overall
Non-negative payoff	0.401***	0.442***	0.405***
	(0.017)	(0.037)	(0.018)
Reporting at t, baseline: truthtelling			
Underreporting	−0.015	0.080	−0.001
	(0.041)	(0.065)	(0.040)
Modest overreporting	−0.027	−0.066	−0.037
	(0.039)	(0.071)	(0.041)
Excessive overreporting	0.097**	−0.090	0.035
	(0.041)	(0.077)	(0.039)
Demographics	✓	✓	✓
Round dummies	✓	✓	✓
Matching group dummies	✓	✓	✓
Observations	600	184	808

Notes: Standard errors in parentheses, $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Marginal effects from pooled probit regressions, with standard errors clustered at the individual level. All specifications include matching group and round dummies and demographic controls (gender, age, geographic origins (with students coming from the central regions of Italy as a baseline), a dummy for students in Economics and the level of experience in lab experiments). Categories of reporting at t correspond to $\frac{\hat{v}-v}{v}$ being either negative (underreporting), null (truthtelling, baseline category), <100% (modest overreporting) or ≥100% (excessive overreporting).

found out lying about the company's value. However, the coefficient measuring whether a lie was detected is small and insignificant.[24]

We summarize these insights in

Result 4:

Sellers are fully opportunistic. They make their acceptance decision dependent on whether the resulting payoff is positive.

5.2.3. Welfare considerations

As anticipated in Section 5.2.1, acceptance occurs significantly more often when information leaks happens. This effect has a clear welfare implication, since trade is always welfare enhancing in the AaC game. Table 6 further analyzes the consequences of v being leaked or not in terms of buyers and seller's payoffs. In particular, we recall the frequency of trades (as reported in Table 3), and complement them with average payoffs, both overall and by participants' roles. Table 6 also shows the differences in the variables between the two information conditions and their significance levels according to $p - values$ from paired t-tests run on matching-group level averages.

Overall, participants' payoff is slightly higher when v is known by both parties, but the difference is not statistically significant. However, when distinguishing between the roles of the participants, a more differentiated picture emerges. As a matter of fact, sellers gain significantly more in the asymmetric information setting (a) while the opposite is true for buyers in setting (b). This result is, of course, not new and descends from the mere exploitation of information advantage. What we find more interesting is the fact that information leaks can not only relieve buyers' disadvantaged position, but also diminish the cases of no trade, which always imply a complete loss of surplus.

[24] The demographic controls are not significantly associated with the probability of acceptance. Coefficients for demographics and matching group and round dummies are available upon request.

Table 6
Summary statistics of trade and payoffs.

	Trade	Payoff			N
		Both	Sellers	Buyers	
(a) v not leaked	73.17%	10.42	12.41	8.43	600
(b) v leaked	80.77%	11.05	10.34	11.75	208
$b - a$	7.6**	0.63	−2.07***	3.32**	

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$, significance level from paired t test run on matching-group level averages ($N = 17$).

This evidence confirms *WITL* and is summarized by

Result 5:

Information leaks increase the probability of trade and, therefore, are welfare increasing.

6. Conclusion

Laboratory experiments exploring (im)moral behavior partly suffer from implicit and explicit demand effects. We tried to weaken this by not only employing a market setup, but also rendering the demand effect for lying ambiguous: although sellers might try to trigger higher price offers via exaggerating value messages, they should be discouraged by leaking information. In this way, information leaks modify the Acquiring-a-Company setup: by default, there is asymmetric value information of the seller who can send a true or false value message. This we have enriched by including a commonly anticipated probabilistic leak-event whose probability is either 10 % or 40%.[25] Via info leaks also the buyer can know the true company value. Thus, the buyer can identify an untrue value message as a lie and also determine its extent and sign before proposing a price which the seller may finally reject or accept. Only in the latter case, trade occurs.

The higher leak probability surprisingly only mildly enhances truthtelling. Instead, it reduces the size of lies and changes the direction of misreporting. Incentives for deviating from the truth are twofold. On the one hand, sellers frequently overstate the value of their company in order to induce higher price proposals and, thus, boost their own payoffs. This overreporting accounts for roughly two thirds of all misreports and decreases when the leak probability increases. On the other hand, we also observe underreporting, possibly due to hoping that it increases the probability of trade. Actually cases of underreporting increase when information leaks are more likely.

Buyers who are unaware of the company's actual value base their price proposal on the value message, although cheap talk. Instead, buyers who found out the actual value of the company anchor solely on the true value and abstain from sanctioning against liars. Furthermore, buyers exhibit a certain inertia in suspicion: having interacted with an overreporting seller makes them decrease their price offers in later rounds. Sellers are fully opportunistic with the sign of the potential payoff being the main driver of their acceptance behavior. As often occurs in market setups, ethical concerns hardly matter. Nevertheless our data indicate that probabilistic leaks foster trade and are, therefore, welfare enhancing. We have also shown how they can have an effect on misreporting, even when their likelihood is low.

From a policy perspective, creating such information leaks, e.g. via random monitoring or audits, should be encouraged. Nevertheless, dealing with private information will always prevail and cannot be neglected. There will be credence goods, insiders on financial markets, hidden information of agents in corporate business, etc. So investigating whereas information leaks may limit the exploitation of the

[25] More generally, commonly anticipated "leak" events allow for continuous classes of hybrid games, in our case including the deterministic border cases where the leak probability is either 1 or 0. See also (Fischer et al., 2021), who theoretically analyze leaks in sequential auctions.

uninformed by better informed parties will remain important, irrespective of whether such practices are legally and socially encouraged. This is where our paper tries to add to the literature.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The data and code for replicating the results of this paper are available at https://osf.io/4n8hy/. All files are licensed under a Creative Commons Attribution 4.0 International (CCBY4.0) license.

Acknowledgments

We thank Sven A. Simon, the participants of the sixth meeting of the Behavioral and Experimental Economics Network (BEEN) in Bologna and of the 2022 Innsbruck Winter Summit on (Un)Ethical Behavior in Markets as well as the audience of the 2022 European ESA Meeting for some useful comments and fruitful discussions. We acknowledge the constructive feedback received from the handling Associate Editor and two anonymous referees. We are grateful to the Max Planck Institute for Research on Collective Goods in Bonn for financing this research project. Tim Lohse is grateful for hospitality of LUISS Guido Carli. The usual caveats apply.

Appendix. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.socec.2023.102151.

References

Abbink, K., & Wu, K. (2017). Reward self-reporting to deter corruption: An experiment on mitigating collusive bribery. *Journal of Economic Behaviour and Organization*, *133*, 256–272.

Abeler, J., Nosenzo, D., & Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, *87*(4), 1115–1153.

Akerlof, G. A. (1978). The market for "lemons": Quality uncertainty and the market mechanism. In *Uncertainty in economics* (pp. 235–251). Elsevier.

Allingham, M. G., & Sandmo, A. (1972). Income tax evasion: A theoretical analysis. *Journal of Public Economics*, *1*(3–4), 323–338.

Alm, J., & Malézieux, A. (2021). 40 Years of tax evasion games: a meta-analysis. *Experimental Economics*, *24*, 699–750.

Anbarcı, N., Feltovich, N., & Gürdal, M. Y. (2015). Lying about the price? Ultimatum bargaining with messages and imperfectly observed offers. *Journal of Economic Behaviour and Organization*, *116*, 346–360.

Angelovski, A., Di Cagno, D., Güth, W., & Marazzi, F. (2020). Telling the other what one knows? Strategic lying in a modified acquiring-a-company experiment with two-sided private information. *Theory and Decision*, *88*(1), 97–119.

Barron, K. (2019). Lying to appear honest. *Technical report*, WZB Discussion Paper.

Barron, K., Kajackaite, A., & Saccardo, S. (2022). Image concerns and lying behavior. Available at SSRN 4204681.

Bašić, Z., & Quercia, S. (2022). The influence of self and social image concerns on lying. *Games and Economic Behavior*, *133*, 162–169.

Bazerman, M. H., & Samuelson, W. F. (1983). I won the auction but don't want the prize. *Journal of Conflict Resolution*, *27*(4), 618–634.

Bénabou, R., & Tirole, J. (2006). Incentives and prosocial behavior. *American Economic Review*, *96*(5), 1652–1678.

Besancenot, D., Dubart, D., & Vranceanu, R. (2013). The value of lies in an ultimatum game with imperfect information. *Journal of Economic Behaviour and Organization*, *93*, 239–247.

Bigoni, M., Fridolfsson, S.-O., Le Coq, C., & Spagnolo, G. (2012). Fines, leniency, and rewards in antitrust. *Rand Journal of Economics*, *43*(2), 368–390.

Bigoni, M., Fridolfsson, S.-O., Le Coq, C., & Spagnolo, G. (2015). Trust, leniency, and deterrence. *The Journal of Law, Economics, and Organization*, *31*(4), 663–689.

Boles, T. L., Croson, R. T., & Murnighan, J. K. (2000). Deception and retribution in repeated ultimatum bargaining. *Organizational Behavior and Human Decision Processes*, *83*(2), 235–259.

Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, *91*(1), 166–193.

Bursztyn, L., Egorov, G., & Fiorin, S. (2020). From extreme to mainstream: The erosion of social norms. *American Economic Review*, *110*(11), 3522–3548.

Buso, I. M., De Caprariis, S., Di Cagno, D., Ferrari, L., Larocca, V., Marazzi, F., Panaccione, L., & Spadoni, L. (2020). The effects of COVID-19 lockdown on fairness and cooperation: Evidence from a lablike experiment. *Economics Letters*, *196*, Article 109577.

Butler, J. V., Serra, D., & Spagnolo, G. (2020). Motivating whistleblowers. *Management Science*, *66*(2), 605–621.

Chavanne, D., & Ferreira, M. (2017). When knowledge is not power: asymmetric information, probabilistic deceit detection and threats in ultimatum bargaining. *Journal of Behavioral and Experimental Economics*, *69*, 4–17.

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

Choshen-Hillel, S., Shaw, A., & Caruso, E. M. (2020). Lying to appear honest. *Journal of Experimental Psychology: General*, *149*(9), 1719.

Croson, R., Boles, T., & Murnighan, J. K. (2003). Cheap talk in bargaining experiments: lying and threats in ultimatum games. *Journal of Economic Behaviour and Organization*, *51*(2), 143–159.

Dana, J., Weber, R. A., & Kuang, J. X. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, *33*(1), 67–80.

Dato, S., Feess, E., & Nieken, P. (2019). Lying and reciprocity. *Games and Economic Behavior*, *118*, 193–218.

Di Cagno, D., Galliera, A., Güth, W., Pace, N., & Panaccione, L. (2017). Experience and gender effects in acquisition experiment with value messages. *Small Business Economics*, *48*, 71–97.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., & Wagner, G. G. (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, *9*(3), 522–550.

Dwenger, N., & Lohse, T. (2019). Do individuals successfully cover up their lies? Evidence from a compliance experiment. *Journal of Economic Psychology*, *71*, 74–87.

Fehr, E., & Gächter, S. (2002). Altruistic punishment in humans. *Nature*, *415*(6868), 137–140.

Fehr, E., & List, J. A. (2004). The hidden costs and returns of incentives—trust and trustworthiness among CEOs. *Journal of the European Economic Association*, *2*(5), 743–771.

Fischbacher, U., & Föllmi-Heusi, F. (2013). Lies in disguise—an experimental study on cheating. *Journal of the European Economic Association*, *11*(3), 525–547.

Fischer, S., Güth, W., Kaplan, T. R., & Zultan, R. (2021). Auctions with leaks about early bids: Analysis and experimental behavior. *Economic Inquiry*, *59*(2), 722–739.

Fries, T., Gneezy, U., Kajackaite, A., & Parra, D. (2021). Observability and lying. *Journal of Economic Behaviour and Organization*, *189*, 132–149.

Geraldes, D., Heinicke, F., & Kim, D. G. (2021). Big and small lies. *Journal of Behavioral and Experimental Economics*, *91*, Article 101666.

Gneezy, U., Kajackaite, A., & Sobel, J. (2018). Lying aversion and the size of the lie. *American Economic Review*, *108*(2), 419–453.

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association*, *1*(1), 114–125.

Güth, W., Marazzi, F., & Panaccione, L. (2023). Exploiting ultimatum power when responders are better informed- theoretical and experimental analysis of conflict resolution. *Journal of Conflict Resolution*, 00220027231167846.

Khalmetski, K., & Sliwka, D. (2019). Disguising lies—Image concerns and partial lying in cheating games. *American Economic Journal: Microeconomics*, *11*(4), 79–110.

Kim, J.-Y. (1996). Cheap talk and reputation in repeated pretrial negotiation. *Rand Journal of Economics*, 787–802.

Kocher, M. G., Schudy, S., & Spantig, L. (2018). I lie? We lie! why? Experimental evidence on a dishonesty shift in groups. *Management Science*, *64*(9), 3995–4008.

Koning, L., Steinel, W., van Beest, I., & van Dijk, E. (2011). Power and deception in ultimatum bargaining. *Organizational Behavior and Human Decision Processes*, *115*(1), 35–42.

Kriss, P. H., Nagel, R., & Weber, R. A. (2013). Implicit vs. explicit deception in ultimatum games with incomplete information. *Journal of Economic Behaviour and Organization*, *93*, 337–346.

Lohse, T., & Simon, S. A. (2021). Compliance in teams–implications of joint decisions and shared consequences. *Journal of Behavioral and Experimental Economics*, *94*, Article 101745.

Meub, L., & Proeger, T. E. (2015). Anchoring in social context. *Journal of Behavioral and Experimental Economics*, *55*, 29–39.

Mitzkewitz, M., & Nagel, R. (1993). Experimental results on ultimatum games with incomplete information. *International Journal of Game Theory*, *22*(2), 171–198.

Reuben, E., & Stephenson, M. (2013). Nobody likes a rat: On the willingness to report lies and the consequences thereof. *Journal of Economic Behaviour and Organization*, *93*, 384–391.

Roth, A. E., & Malouf, M. W. (1979). Game-theoretic models and the role of information in bargaining. *Psychological Review*, *86*(6), 574.

Samuelson, W. F., & Bazerman, M. H. (1985). Negotiation under the winner's curse. *Research in Experimental Economics*, *3*, 105–138.

Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, *4*, 25–55.

Selten, R., Sadrieh, A., & Abbink, K. (1999). Money does not induce risk neutral behavior, but binary lotteries do even worse. *Theory and Decision*, *46*(3), 213–252.

Serra, D. (2012). Combining top-down and bottom-up accountability: evidence from a bribery experiment. *The Journal of Law, Economics, & Organization*, *28*(3), 569–587.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *Science*, *185*(4157), 1124–1131.

Utikal, V., & Fischbacher, U. (2013). Disadvantageous lies in individual decisions. *Journal of Economic Behaviour and Organization*, *85*, 108–111.

Vesely, S. (2014). Ultimatum game with asymmetric information: A study of deception and fairness. *Studia Psychologica*, *56*(1), 53.