NBER WORKING PAPER SERIES

ETHICS AND ILLUSIONS:
HOW ETHICAL DECLARATIONS SHAPE MARKET BEHAVIOR

John M. Barrios
Jeremy Bertomeu
Radhika Lunawat
Ibrahima Sall

Working Paper 32385
http://www.nber.org/papers/w32385

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2024

We are grateful to seminar participants at HKUST, Odense, Washington University in St Louis, Carnegie Mellon, and American University for suggestions on earlier versions of the paper. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Ethics and Illusions: How Ethical Declarations Shape Market Behavior
John M. Barrios, Jeremy Bertomeu, Radhika Lunawat, and Ibrahima Sall
NBER Working Paper No. 32385
April 2024
JEL No. D53,G10,G14,G4,G41

ABSTRACT

We examine the impacts of ethical declarations on market transactions through a controlled laboratory experiment, where privately-informed sellers issue a public report prior to a first-price auction. We find that while signing an ethical statement does not reduce misreporting by sellers, it significantly increases buyer trust, often skewing the terms of the trade in favor of sellers. Contrary to rational expectations, buyers consistently struggle to undo the bias. In counterfactual scenarios, from our structural analysis, we find that price efficiency improves when buyers rationally process uncertainty about sellers' ethical preferences, yet bias persists even when buyers have more accurate perceptions of sellers'' ethical standards. Overall, our results suggests that disclosure interventions aimed at enhancing ethical conduct in market settings may not necessarily lead to more efficient pricing or reduced bias, and in some instances, may even disadvantage certain market participants.

John M. Barrios
Olin Business School
Washington University in St. Louis
1 Brookings Drive
St. Louis, MO 63130
and NBER
John.barrios@wustl.edu

Jeremy Bertomeu
Washington University in St. Louis
jeremy.bertomeu@gmail.com

Radhika Lunawat
University of California, Irvine
rlunawat@uci.edu

Ibrahima Sall
Washington University in St. Louis
i.sall@wustl.edu

"To get markets that work, you have to keep the other person from trying to cheat you at every moment. So morality is closely related to the workings of the market."

The Economy of Trust: An Interview with Kenneth Arrow (Arrow 2006)

1 Introduction

The dynamics of market transactions are not merely dictated by supply, demand, and regulatory frameworks, but also profoundly influenced by the ethical conduct and disclosure practices of market participants (Smith 1776; Stevens 2023). Given that not all statements can be verified, honesty in communication is central for the effective functioning of society, ensuring integrity and reliability in various sectors, from governmental agencies to scientific research (Marco and Larkin 2000; Yang 2009; Kristal et al. 2020). The role of ethics in economics, particularly in financial markets, goes beyond mere compliance with legal standards; it encompasses a broader spectrum of moral principles and transparency that underpin trust and integrity in market interactions (Arrow 2006). In this context, the explicit disclosure of ethical stances, or the lack thereof, plays a pivotal role in shaping market behaviors and expectations.

Despite the recognized importance of ethics in markets, theoretical models often inadequately capture its complex impact on trade. Traditional theoretical research has shown that, if the benefits of misreporting are common knowledge, market participants should rationally anticipate the bias and, in equilibrium, correct the report to form beliefs based on the true information. This theoretical expectation is grounded in the concept of a rational expectations equilibrium, where no individual is misled by bias (Dye 1988; Stein 1989). This assumption, however, falters in the real world, where many common examples appear to reflect predictable bias, such as exaggerating income or height in a personal exchange, managing earnings upwards to meet a target, embellishing a product in an ad, exaggerating achievements in a vita, among many other examples.

Our study seeks to bridge the existing gap in understanding how statements about ethical behavior influence market exchange. We utilize a market experiment where a privately informed

seller auctions an asset to competing buyers and communication is potentially distorted by the costs associated with misreporting. We address three primary research questions: (1) Given an understanding of the seller's economic incentives, can buyers undo the seller's reporting bias? (2) To what extent does heterogeneity in aversion to lying contribute to noise in communication? (3) How does an intervention, which requires the seller to sign an ethical statement, impact the terms of trade and the overall quality of communication?

These questions are deeply rooted in the foundational principles of well-functioning markets, tracing back to Adam Smith's 1759 Theory of Moral Sentiment (see, section I, *"The prudent man is always sincere, and feels horror at the very thought of exposing himself to the disgrace which attends upon the detection of falsehood"*). This perspective is further echoed in contemporary analyses of trust problems such as the tragedy of the commons and the prisoner's dilemma (Kreps et al. 1982) as well as in the context of measuring social capital or investor trust (Barrios et al. 2021b; Guiso et al. 2008a). Our argument builds upon these foundations, proposing that the interplay between ethics, trust, and trade is more complex than commonly perceived. The heterogeneity of ethical preferences among individuals, coupled with the inherent challenges in making direct observations, generates a phenomenon we label as "ethical noise." Ethical noise confounds the process of accurately inferring predictable bias, standing in stark contrast to the prediction of perfect communication under rational profit-maximization (Dye 1988; Stein 1989).

Our empirical approach employs a laboratory implementation of the classic single-period model of strategic communication, which incorporates costs for lying (Guttman et al. 2006; Kartik 2009; Kartik et al. 2007). This model has been extensively referenced in the literature exploring strategic communication and misreporting (Dye 1988; Stein 1989; Einhorn and Ziv 2008; Bertomeu et al. 2013; Caskey and Laux 2013). In our experimental framework, a seller, having privately assessed the value of an asset, makes a public report prior to sale. The cost associated with reporting escalates in proportion to the degree of bias in the information provided. Additionally, in separate treatment conditions, sellers may be mandated to sign a public ethical statement before proceeding with the sale, or they might have the option to either sign or abstain from signing such a statement.

The monetary payoffs for the sellers are contingent upon the extent of the bias in their report and the final winning bid. Buyers, in turn, observe these public reports and are required to estimate the true value of the asset, subsequently placing bids in a first-price auction. For buyers, the monetary rewards are linked to the accuracy of their value prediction and the net profit from winning the auction. The roles of buyers and sellers are rotated in each new round, with participants being re-matched randomly across the experimental pool in groups of five.

We document several experimental results. First, buyers were generally unable to correct for the reporting bias of sellers. The average error in winning bidders' individual prediction was about \$17.53 (experimental dollars), while the asset's average fundamental value was approximately \$48.70 across all rounds of the experiment. Interestingly, when considering the collective average estimate of all buyers, including those not winning the auction, the deviation from the true value was small (roughly \$1.77, measured as the average difference between estimated and actual asset values). This discrepancy suggests that while individual beliefs were often skewed, the collective wisdom of the crowd seemed to provide a more accurate and unbiased estimate. This phenomenon points to the crowd's ability to average out *strategic* errors, which, in our experiment, did not stem from the aggregation of private signals about the true value (van Dolder and van den Assem 2018; Da and Huang 2020).

This finding is in stark contrast to the predictions of rational expectations with strict (monetary) payoff maximization, where buyers, understanding the equilibrium strategy of the seller, are expected to correctly infer the private information and consequently bid the true value of the asset netting out from the report the equilibrium bias. Moreover, in the baseline treatment, variability in the individual prediction errors translated into heterogeneity in bidding behavior and more importantly winning buyers overbid by roughly \$15 on average; thus, sellers profited from bias.

Second, we observe the varied inference processes resulting from differing seller preferences for honesty. Echoing prior literature (Evans et al. 2001; Mittendorf 2006; Hurkens and Kartik 2009; Gneezy et al. 2013; Abdel-Rahim and Stevens 2018; Gneezy et al. 2018), we find a bimodal distribution of ethical preferences among sellers. This distribution features a peak at truthful reporting

and a (smaller) peak at payoff maximization. Such a distribution pattern is indicative of the uncertainty bidders face regarding the forthcoming nature of the seller (Marinovic 2013; Aghamolla and Smith 2022; Jiang et al. 2022). We conduct the same exercise to analyze the beliefs of buyers based on their estimated biases. We find a more diffuse distribution. This diversity in beliefs translated to a disproportionate representation of buyers who assumed seller honesty among the auction winners. In a counterfactual analysis assuming a rational expectations equilibrium framework, where buyers accurately infer the distribution of seller preferences and sellers act accordingly, ethical noise would have resulted in average inference errors equal to \$6.3 compared to the \$8.9 obtained from the experiment. Even when accounting for the change in reporting strategies given common knowledge rationality, our evidence thus suggests that rational play increases communication.

Next, we focus on the impact of making salient the role of ethics in the auction by augmenting our experimental design with treatments to assess the impact of ethical codes on market behavior (Davidson and Stevens 2013). In these treatments, sellers were either required to (a) sign an ethical statement before participating in the auction ("mandatory") or (b) presented with a choice to voluntarily sign or decline such a statement ("voluntary") with buyers observing the decision of the seller. Contrary to our expectation, the introduction of ethical statements did not significantly alter the reporting behavior of sellers, irrespective of the statements.[1] However, despite the impact on sellers being small, we found a significant effect on buyers' behavior. The ethical statement was perceived by buyers to imply that sellers were more honest, which led to greater forecast errors and overbidding. For example, in scenarios with a signed ethical statement, the average absolute error in buyers' forecasts was \$8.95 and their bids were on average \$7.66 higher, resulting in buyers being worse-off by \$4.54. This outcome indicates that ethical interventions may have a more pronounced effect on buyers, potentially leading to increased mispricing.

[1] We recognize that this could suggest that our experimental setup might not have had sufficient power to induce notable shifts in honesty. Our aim was not to design a mechanism to enhance ethics but rather to observe the implications of a reasonably framed ethical statement on buyers and sellers, and we used for this purpose an ethical statements from prior literature. Previous research (Mazar et al. 2008; Davidson and Stevens 2013; Abdel-Rahim and Stevens 2018) underscores the efficacy of certain interventions in influencing honesty. Further, while the effect on sellers was insignificant, our market design allows us to find a significant effect on buyers, suggesting that the effect of ethics intervention may truly be greater on buyers and not simply reflecting a lack of power in the design.

Our study offers several significant contributions to the understanding of ethics and trust in capital markets. An extensive literature highlights the pivotal role that trust plays in supporting well-functioning markets, especially when contracts are either incomplete or unenforceable. Motivating our study is the work of Sapienza et al. (2013), who show that global surveys on trust often measure perceptions of others' honesty rather than individuals' own reporting behaviors. Our experiment uniquely allows us to disentangles these two aspects - beliefs about trust and actual honesty - within the same laboratory market setting.

Moreover, our study is closely aligned with a substantial body of work in institutional economics that illustrates how institutions fostering trust enhance economic exchange, see, e.g., Milgrom et al. (1990), Guiso et al. (2008b), Sapienza and Zingales (2012), Acemoglu et al. (2020), or Barrios et al. (2021a). Three particular studies in this area share a thematic connection with ours, focusing on the credibility of reports when potential misalignment exist. Jin and Kato (2006), in the context of online baseball card auctions, observe that many sellers readily accept high-quality claims at face value. Elfenbein et al. (2019) examine eBay charity listings, noting a tendency for unverified charity auctions to fetch high prices despite verification typically increasing the likelihood of a sale. Sheremeta and Shields (2013) explore a cheap talk experiment and find a tendency for subjects to expect more honesty from others than they exhibited themselves in similar roles. While these studies focus on the interpretation of cheap talk, they suggest a prevailing expectation of honesty among receivers, potentially more than senders. Taken together with our results, this raises a question about the effectiveness of disclosure interventions aimed at promoting honesty. Such interventions might inadvertently heighten receiver beliefs more than they influence actual honesty.[2]

Additionally, our study contributes to the understanding of trust in a two-sided market setting

[2]Our study gains additional relevance in light of recent developments in the field of ethical research. Notably, the retraction of an ethical signatures study published in the Proceedings of the National Academy of Sciences (Kristal et al. 2020) which has highlighted the complexities and challenges inherent in reliably measuring and interpreting ethical behavior in experimental settings. This development underscores the necessity for rigorous and transparent research methodologies, particularly in the delicate intersection of ethics and market behavior. Our study, with its robust experimental design and comprehensive approach, aims to contribute constructively to this ongoing discourse, providing insights reflective of the nuanced nature of ethics in economic transactions.

involving transactions between buyers and sellers (O'Brien and Srivastava 1991; Forsythe et al. 1999; Bloomfield and O'Hara 1999; Crockett et al. 2009; Bloomfield et al. 2009). Commonly, trust is measured either through survey questions like "Do you believe most people can be trusted or should you be careful in dealing with people?" or through laboratory trust games where one player can increase or decrease the value of another player's investment (Berg et al. 1995; Glaeser et al. 2000). Our experiment diverges conceptually in two key ways. First, in our market setup, participants are dividing the gains from trade, unlike trust games where players engage in joint production with mutual cooperation motives. Additionally, previous research indicates that market-mediated exchanges might actually erode moral sentiment (Falk and Szech 2013). Second, trust pertains to biases in reporting about a fundamental value, while in trust games with communication, it often relates to the credibility of promises about future cooperation. In summary, the nature of our experiment aims to capture features salient in pre-trade communication.

Our research builds upon studies that have explored the failures of rational expectations in scenarios involving verifiable disclosure, where a privately informed sender may choose to either truthfully disclose their information or withhold it (Dickhaut et al. 2003; Bourveau et al. 2020; Jin et al. 2021, 2022). While rational expectations theory posits that such scenarios should unravel to complete transparency (Milgrom 1981), empirical evidence often shows a struggle to achieve full information disclosure; however, theoretically, failure of unravelling is usually explained in terms of frictions to communication such as non-verifiability or costs (Beyer et al. 2010). Our setting involves a strategic communication game that demands high-level strategic thinking, posing significant challenges in processing information based on other players' actions (Corgnet et al. 2018) and diverging from the verifiable disclosure literature on several conceptual dimensions. For instance, lying by omission, as in withholding information in a disclosure setting, nay differ from misreporting. In the former, the sender makes a choice that is fully recognized by the receiver, without necessarily a moral expectation of full disclosure. This nuance means that there is potentially less scope for ethical triggers compared to untruthful reporting. Furthermore, the nature of the sender's actions differs: under costly misreporting, the sender is not deceiving the receiver

in equilibrium as the true signal can be discerned (Dye 1988; Stein 1989), whereas under truthful disclosure, withholding reduces the overall informativeness of the communication.

Our analysis is organized along four sections. Section 2 presents the theoretical framework and empirical design, including the information known by buyers and sellers and the market structure in the experiment. Section 3 provides both descriptive and reduced-form evidence on the relation between ethics and misreporting. Section 4 focuses on implementing estimates of ethical behaviors using a formal model of misreporting, evaluated from the viewpoints of both buyers and sellers. We conclude the paper in Section 5.

2 Theoretical Framework and Experimental Design

2.1 Theoretical Framework

Our experiment builds upon established misreporting models with noisy misreporting á la Fischer and Verrecchia (2000) and Frankel and Kartik (2019). These models provide a foundation for understanding sellers' behavior in the presence of informational asymmetries and allow us to examine ethical considerations.

In the model, each seller observes the true value of the asset, denoted as $\tilde{v}$. This value follows a cumulative distribution function (c.d.f.) $F(.)$, which is uniformly distributed over the interval $[0, 100]$. The seller also possesses an intrinsic attribute referred to as "ethics type," captured as $\tilde{e}$ with c.d.f. $G(.)$. This type influences the seller's decision-making process, particularly in how they balance honesty against profit maximization. The realized type e quantifies this balance: $e < 1$ indicates a leaning towards ethical behavior, valuing honesty more than personal gain, whereas $e = 1$ captures a seller who is solely focused on maximizing their expected payoff with no regard for ethical considerations.

After observing the true value, the seller makes a public report r to maximize a selling price

net of misreporting costs.[3] This choice is captured by the following optimization problem:

$$R(v,e) \in \underset{r}{\operatorname{argmax}}\, \pi(r|v,e) \equiv eP(r) - K(r,v), \tag{1}$$

where $K(r,v) = (r-v)^2/20$ is an increasing positive function and $P(.)$ is pricing function which the seller takes as given.

There is a market with well-diversified investors pricing the asset at its expected value conditional on the public report. Investors use Bayes' rule to update their beliefs about the true value of the asset, $\tilde{v}$, based on the seller's report. Therefore, for any report r such that Bayes rule applies, the price satisfies:

$$P(r) = \mathbb{E}(\tilde{v}|r = R(\tilde{v},\tilde{e})). \tag{2}$$

This provides us with our definition of market equilibrium.

Definition 1 *An equilibrium is defined by a reporting strategy R^* and a pricing strategy P^* such that the report R^* solves (1) evaluated at $P = P^*$ and P^* satisfies (2) evaluated at $R = R^*$.*

Our analysis focuses on equilibria characterized by differentiable strategies and pricing functions, where higher reports are consistently interpreted by investors as indicating more favorable asset values, that is, $P'(r) > 0$. When the seller's problem yields an interior solution, the optimal report r meets the first-order condition in equation (1) is

$$eP'(r) - K_r(r,v) = 0. \tag{3}$$

As a benchmark, consider the special case where it is common-knowledge that sellers are primarily interested in maximizing revenue, i.e., setting $e = 1$. In a separating equilibrium, the reporting bias is given by $r - v = 10P'(r)$. This results in a pricing function given by $P(r) = r - 10P'(r)$ such that all buyers can undo the bias from their knowledge of the equilibrium pricing function. As noted by Einhorn and Ziv (2012), the equilibrium requires that the lowest type (the

[3] If e is positive, one can equivalently interpret the parameter as a benefit from reporting or a random cost, as both sides of (1) can be divided by $1/e$ so that the implied cost is simply $K(r,v)/e$. However, the model with random benefits is more flexible because it allows us to interpret sellers with negative e as a preference to under-report which, holding prices fixed, tends to protect buyers against overbidding.

seller observing $v = 0$) attains no less than zero (and therefore does not bias), that is, $P(0) = 0$, implying a pricing function

$$P(r) = r - 10(1 - e^{-r/10}). \tag{4}$$

The bias, $10(1 - e^{-r/10})$ starts from zero for the lowest report and progressively increases toward a maximum bias of 10 as the reports increase.[4]

In our experimental setup as discussed below, we are primarily interested in inferring two key distributions: the *actual* distribution of ethical preferences as revealed by the seller's reports and the *perceived* distribution of these ethical preferences, as inferred from the buyer's estimates. To do so, it is convenient to invert the reporting strategy $R(v, e) = r$ in equation (3) so that

$$e = K_r(r, v)/P'(r), \tag{5}$$

where e captures the implied ethical disposition of a seller who reports a value r for an asset with a true value v. This approach follows the empirical strategy for estimating ethics distribution, as suggested by Bertomeu et al. (2019). To structurally estimate the actual ethics e in our experiment, we can use the formula $\hat{e} = K_r(r, v)/\hat{P}'(r)$, where $\hat{P}'$ is the derivative of the estimated selling price. Correspondingly, perceived ethics is calculated using $\hat{e}^E = K_r(r, v^E)/\hat{P}'(r)$, where v^E stands for the buyer's estimate of the seller's true asset value.

2.2 Experimental Design

Having established the theoretical underpinnings of our model, particularly the mechanisms through which sellers' ethical preferences impact their reporting behavior, we now turn our attention to the practical implementation of these concepts in the experimental design. The design aims to empirically test the theoretical model, offering insights into how actual and perceived ethical behaviors manifest in a controlled market environment.

[4] See Einhorn and Ziv (2012) for a general solution to this problem; the boundary condition $P(0) = 0$ is required from the fact that, in a separating equilibrium, a seller with value 0 would be better-off reporting truthfully over any other report. As noted in Guttman et al. (2006) and Guttman et al. (2010), this game also has ex-ante preferred semi-pooling equilibria in which some (off-equilibrium) messages are not be played. Because we did not guide subjects toward these equilibria, we did not find, in the experiment, evidence of these equilibria and the estimated price function is consistent with differentiability.

Our experimental sessions consist of several independent rounds, each comprising key components detailed in Appendix A. These components include displays presented to subjects (from windows A.0-A.5), a pre-experiment questionnaire, and practice sessions designed to familiarize participants with the game.[5]

In each period of the experiment, participants are randomly assigned roles as either a seller, who is endowed with an asset, or as one of four buyers, each endowed with 100 experimental dollars. Prior to the start of the experiment, each subject receives a full description of experimental procedures. They are asked to answer six questions about the rules of the experiment (window A.0), to ensure an understanding of the experimental structure. After this, participants engage in five practice rounds. These rounds are used to ensure that subjects understand the game, but do not contribute to the calculation of final payoffs. The core of the experiment consists of 20 anonymous paying rounds, grouped randomly for each round, drawn from a pool of 15 to 25 subjects to prevent the formation of reputational biases.[6] The anonymity of participants is maintained throughout each round to eliminate any reputational concerns that might add another consideration to decision-making and are not our current focus. We adopt a pay-one approach to mitigate the effect of dynamic wealth effects and ensure incentive-compatibility (Charness et al. 2016; Azrieli et al. 2018). A single paying round is drawn randomly and a participant receives a base payment of \$7, plus any net profit earned in the selected round calculated at a rate of 16.67 cents per experimental dollar.

During each round, the value v of an asset is independently drawn from a uniform distribution ranging from $\{0, 1, \ldots, 100\}$. The seller, upon observing this draw, chooses a report r, which is then communicated with all bidders. Notably, the report r is not restricted to whole numbers and can fall outside the $[0, 100]$ (window A.1), which we allow because, in the equilibrium in (4), a

[5] A demo of the experiment is available at http://stag-klapp.herokuapp.com/demo for further reference.

[6] It would be interesting to examine the dynamic formation of reputations in disclosure, which involves additional considerations as to how credible reputations might form in the presence of an intervention on ethics. Our research question, however, is primarily targeted at decentralized exchanges where not all senders could form reputations. We made this design choice, not because reputations are unimportant, but because they add a supplementary mechanism that would likely make it much more difficult to theoretically separate the relative contribution of reputations and ethics, and contribute additional questions as to whether the mechanisms are complements or substitutes. We leave this important question for further research.

fully-revealing report may require types with $v \geq 90$ to send reports above 100. The reporting range is centered at the true value, permitting reports of the true value plus or minus 50, since any amount beyond this range would lead to prohibitive misreporting costs.

In the "ethical" treatments, sellers face a choice regarding the certification of their report. They are asked whether they want to certify the report and, in the affirmative, sign electronically a disclosure form stating that the report is truthful (window A.2). This ethical certification, however, is effectively cheap talk in the experiment; it does not legally bind the seller or limit their reporting options. We explore three distinct treatments in this regard: (1) no ethics statement, (2) mandatory ethics statement, and (3) a voluntary ethics statement such that the seller may decline to sign. In (2) the mandatory ethics statement scenario, a refusal to sign results in the termination of trading for that session with the seller receiving no payoff and the asset remaining unsold. Buyers are informed about the experimental treatment in play and, in the case of voluntary ethics statements, whether the seller opted to sign or not. The cost of misreporting the true signal is $(r - v)^2/20$ and is borne privately by the seller.

The selling process follows King (1996) and the asset is sold in a first-price auction with four competing bidders. Buyers are tasked with making two decisions in each round: (i) estimating the true value g of the asset, and (ii) submitting a bid b for the asset (window A.3). The auction is won by the highest bidder, with the winning formula for the highest bidder being $100 - |g - v| + (v - b)$, and all other bidders receiving $100 - |g - v|$. In cases where this calculation results in a negative payoff, the net profit for the bidders is set to zero. At the end of each round, buyers learn whether they won the auction, the true value of the asset, and their net profit for that round (window A.4).

At the end of the experiment, as part of the post-experiment procedure, subjects are asked to complete a questionnaire designed to assess their trustworthiness and their trust in others, following the framework established by Davidson and Stevens (2013), in window A.5. These responses are collected in window A.5. Sessions were conducted online, with participants recruited from the subject pool of the Experimental Social Sciences laboratory at a prominent university.

It is worth noting that our experiment's objective is not to engineer systems for enhancing

ethical reporting. For instance, we did not explore the impact of more forceful inducements against dishonesty, alternative framing strategies as suggested by Kristal et al. (2020), or the activation of moral triggers prior to the experiment, a concept explored by Mazar et al. (2008). Instead, our focus was on employing a broad code of ethics, akin to the one used by Davidson and Stevens (2013), to examine how market participants may respond to ethical declarations.

3 Reduced-Form Analysis

3.1 Descriptives

We ran three sessions for each of the three treatments, (i) no ethical statement, (ii) voluntary ethical statement, and (iii) mandatory ethical statement. Figure 1 provides descriptive evidence on participants. Specifically, Figure 1 shows that sellers over-reported on average, buyers' predictions and bids were below the sellers' reports, buyers' bids were below their predictions, and, as is commonly-observed in first-price experiments, winning bids were higher than the seller's report and asset value (Goeree et al. 2002; Crawford and Iriberri 2007).

Figure 1: Mean Seller Reports, Buyer Estimates, and Bids

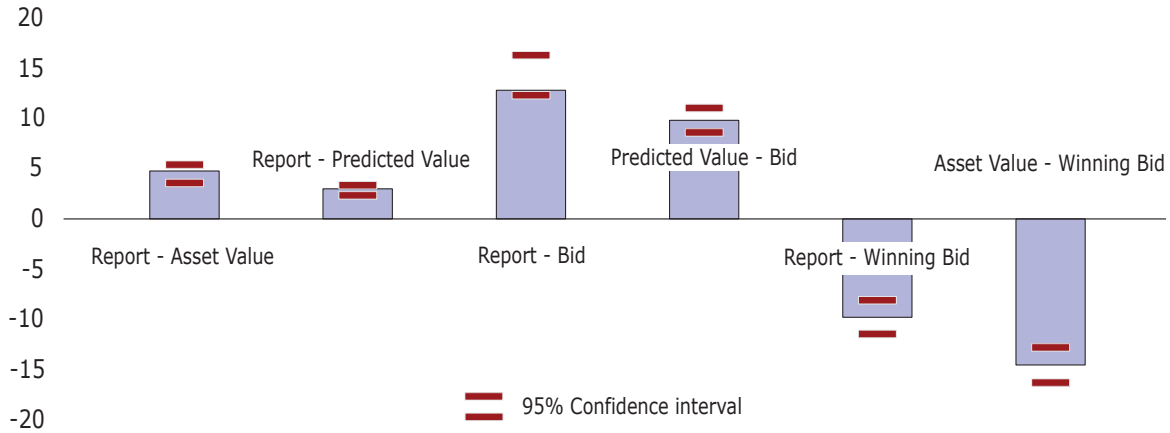

Figure 2 provides a more detailed description of sellers' and bidders' actions across true and reported values. We present bin scatter graphs of sellers' reported values as a function of the asset values in Panel (a) and buyers' bids as a function of reported values in Panel (b), controlling for

age, gender, and risk preferences. In Panel (a), sellers over-reported by $4.95 on average, which is slightly over half of the average bias in the rational expectations solution in equation (4). The reporting bias is roughly constant across asset values, rather than increasing as in the rational expectations equilibrium. We also find evidence that bids appear to take bias into account, as they appear lower than reported values by roughly $13.11 on average. Contrary to actual seller behavior, buyers exert an increasing discount for higher reports, and bid above the report for sufficiently low (i.e. reports below $25).

Figure 2: Sellers Reports and Buyers Bids

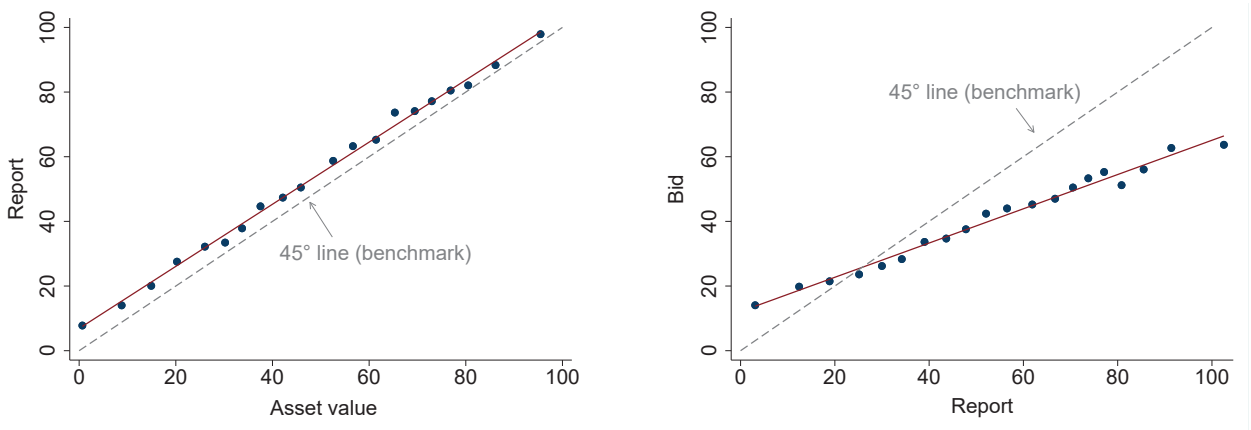

(a) Reported Value as a Function of Asset Value

(b) Bids as a Function of Reported Value

Notes - This figure presents binned scatter plots of sellers' reports as function of asset values in Panel (a) and buyers' bids as function of sellers' reports. The 45-degree° line represents truthful reporting in Panel (a), and buyers bidding the reported value in Panel (b).

3.2 Sellers and Ethics

In Figure 3, we partition the sellers' reports based on the ethical treatment and plot the average reports over the distribution of asset values, controlling for age, gender, and risk preferences. Across all treatments, in Panels (a) and (b), we did not find evidence that signing an ethical statement affected sellers reporting bias, relative to no ethical statement at any asset value.[7] The average

[7]In untabulated analysis we conduct two-by-two mean comparisons of sellers' reporting bias between ethical treatments as well as between each ethical treatment and no treatment, using a t-test. Instead of finding evidence of decreasing bias with the level of ethical treatment, we find no statistical significance in sellers reporting bias mean differences between mandatory and voluntary treatments or between voluntary and no treatment, but higher mean reporting bias in the mandatory ethical treatment setting when compared to the setting without ethical signature.

seller bias is $4.01 under no ethical statement, $5.85 under mandatory signing and $4.86 under voluntary signing. We find more variation for subjects declining to sign, but this may be due to a small sample as fewer sellers declined to sign.

Figure 3: Reported Value as a Function of Asset Value By Ethical Treatment.

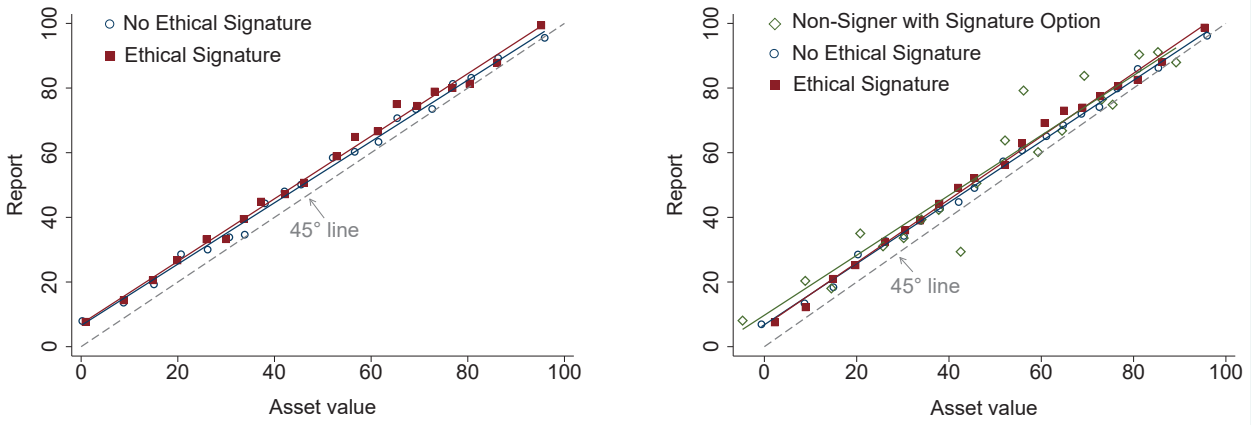

(a) Ethical Signature vs. No Signature

(b) Ethical Signature, Decline to Sign vs No Signature

Notes - This figure presents binned scatter plots of sellers' reports as function of asset values. Panel (a) defines treated sellers as all sellers exposed to an ethical signature (including both sellers who signed the statement under the voluntary arm and those signing when the statement was made mandatory), while in Panel (b) we examine sellers signing the ethical statement in both the voluntary and mandatory treatments as well as sellers declining to sign when signing the ethical statement is made voluntary. The 45-degree° line represents truthful reporting.

To complement Figure 3 and assess statistical significance, we run a multivariate regression with the following following specification:

$$Log\ Report = \beta \times Ethics + \sum \gamma_k \times Control_k + Round\ FE + \epsilon, \tag{6}$$

where $Log\ Report$ is the natural log of the sellers' report plus one, $Ethics$ is an indicator variable for the ethics treatment. In Columns 1-3, this variable is set to one if a seller is presented with an ethical statement to sign. In Columns 4-5, we decompose $Ethics$ into two separate indicators $Signed\ Ethical\ Statement$ and $Declined\ to\ Sign$, equal to one when, in the voluntary treatment, the seller chooses to sign or declines to sign, respectively. $Controls$ is a vector of controls for the individual seller (age, gender, risk preference) as well as the asset value given to the sender. In some specifications, we include fixed effects for each round of play $Round\ FE$.

Results in Table 1 are consistent with Figure 3. In panel A, the ethical treatment combines the voluntary and mandatory signature sessions, while in panel B the voluntary and mandatory ethical treatments are separated. In either case, we fail to find a statistically significant difference in the reporting behavior of sellers exposed to the ethical treatment. This is true when we examine simple exposure to an ethical statement in Columns 1-3 of panel A as well as when we separately examine signing and declining to sign the statements in Columns 4-6 of panel A or Columns 1-6 of panel B. In addition to being statistically insignificant, the differences in reports between sellers exposed to an ethical statement and those who were not exposed are economically small and not with the expected sign: for example, in Column 3 of panel A, the coefficient on $Ethical\ Statement$ of 0.076 translates to sellers exposed to ethical statement reporting nearly \$3.95 *higher* reports on average. This is similar in magnitude to the coefficient from sellers that signed the report in Column 6 of panel A. We also observe similar magnitude when further separating sellers who signed under the mandatory ethical treatment from those signing under the voluntary ethical treatment in Column 6 of panel B. Overall, the reduced form results suggest that the ethical treatment did little to change the reporting behavior of sellers.

Table 1: Regression Analysis of Seller Reports

Panel A	*Log Report*					
	(1)	(2)	(3)	(4)	(5)	(6)
Ethics: Ethical Treatment	0.091	0.095	0.076			
	(0.073)	(0.068)	(0.049)			
Ethics (=1) : Signed Ethical Statement				0.104	0.109	0.073
				(0.073)	(0.068)	(0.049)
Ethics (=2) : Declined to Sign				-0.017	-0.017	0.095
				(0.159)	(0.157)	(0.104)
Observations	677	677	677	677	677	677
Adjusted R-squared	0.001	0.091	0.65	0.001	0.091	0.65
Mean. Report	53.83	53.83	53.83	53.83	53.83	53.83
Controls	No	No	Yes	No	No	Yes
Round FE	No	Yes	Yes	No	Yes	Yes
Panel B	*Log Report*					
	(1)	(2)	(3)	(4)	(5)	(6)
Mandatory : Signed Ethical Statement	0.109	0.11	0.076			
	(0.079)	(0.078)	(0.059)			
Voluntary : Signed Ethical Statement				0.099	0.104	0.083
				(0.086)	(0.079)	(0.054)
Voluntary : Declined to sign				-0.017	-0.012	0.099
				(0.160)	(0.156)	(0.105)
Observations	438	438	438	487	487	487
Adjusted R-squared	0.001	0.081	0.634	0.000	0.096	0.627
Mean. Report	53.95	53.95	53.95	53.05	53.05	53.05
Controls	No	No	Yes	No	No	Yes
Round FE	No	Yes	Yes	No	Yes	Yes

Notes - This table summarizes the OLS regressions results of the amounts reported by the sellers on : ethical treatment (whether voluntary or mandatory), and sellers choices regarding the signature of ethical statement in the voluntary treatment. The baseline is comprised of the amounts reported in the sessions with no ethical treatment. In panel A, the ethical treatment combines the voluntary and mandatory signature sessions, while in panel B the voluntary and mandatory ethical treatments are separated. The regressions are based on the sample with sellers only. *Log Report* is the natural log of one plus the seller's report. *Ethics* is an indicator variable that takes the value of 1 if the seller is presented with an ethical statement to sign (voluntary or mandatory treatment); zero otherwise. *Ethics* is further decomposed into *Signed Ethical Statement* and *Declined to Sign*; indicator variables that take the values of 1 and 2 when the seller signs or declines, respectively; zero otherwise. Controls are comprised of individual seller characteristics (age, gender, risk preference) as well as the true asset value. Round FE indicates whether the estimates are within rounds. The standard errors are clustered at the seller level, and reported in parentheses below the coefficients. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

3.3 Bidders and Ethics

Next, we focus on bidders' behavior, similarly partitioning the bids based on whether the bidders were exposed to the ethical treatment. Figure 4 plots the average bid for treated bidders over the distribution of seller provided reports. Each bin scatter controls for the age, gender, and risk preferences of the participants. As with the sellers in Figure 3, we examine two versions of our ethics treatment in Figure 4. Panel (a) defines treated bidders as all bidders exposed to an ethical signature, while, in panel (b) we separately examine bidders as a function of whether they observe a voluntary signer or a seller declining to sign.

Figure 4: Bids as a Function of Reported Values By Ethical Treatment

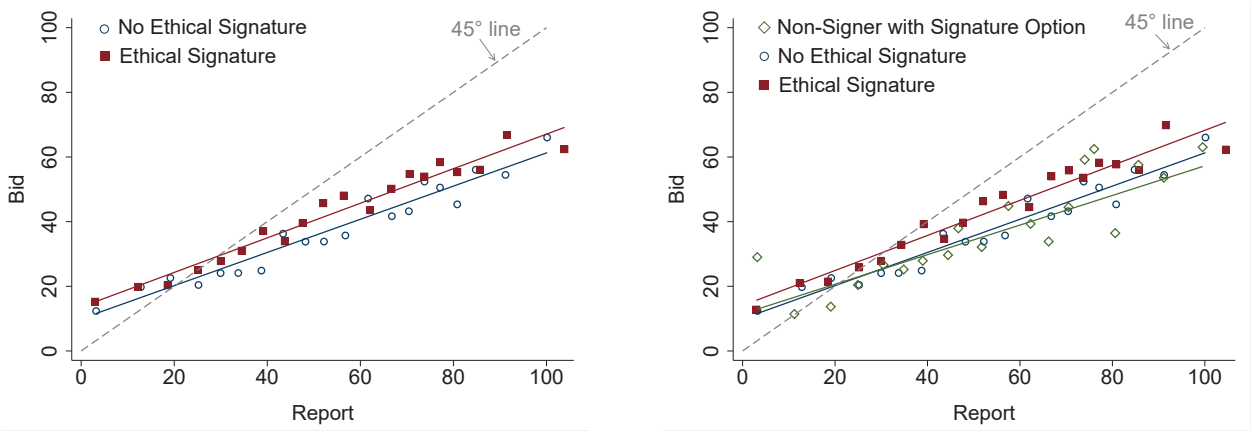

(a) Ethical Signature vs. No Signature

(b) Ethical Signature, Decline to Sign vs No Signature

Notes - This figure presents binned scatter plots of buyers' bids as function of sellers' reports. Panel (a) defines treated buyers as all buyers observing that sellers are exposed to an ethical signature (including both sellers who signed the statement under the voluntary arm and those signing when the statement was made mandatory), while in Panel (b) we examine buyers observing sellers who sign the ethical statement in both the voluntary and mandatory treatments as well as sellers who declining to sign when signing the ethical statement is made voluntary. The 45-degree° line represents buyers bidding the reported value.

Figure 4 suggests that buyers respond to the ethical treatment by choosing higher bids relative to the non-treated buyers. Both panels reveal higher bids from buyers who observe the seller signing an ethical statement. Panel (b) also shows weak evidence that buyers observing sellers declining to sign tend to bid lower as compared to both sellers without an ethical statement and sellers having a sign a statement. The average buyer placed a bid lower than the reported value

by nearly \$15.76 under no ethical statement, \$9.61 under mandatory signing and \$12.36 under voluntary signing.[8] Nonetheless, difference in average bid between settings certainly relate to the amounts reported by sellers.

We then run the following multivariate regression to shed light on bidding behavior when facing a seller who signs an ethical statement, compared to the baseline of no ethical statement treatment.

$$Log\ Bid = \beta \times Ethics + \sum \gamma_k \times Control_k + Round\ FE + \epsilon, \tag{7}$$

where $Log\ Bid$ is the natural log of the bid plus one and all other variables are defined as in equation (6), except that we now include the amount reported by the seller in the control variables, instead of the asset value. Table 2 presents the results from the regression in equation (7). In panel A, the ethical treatment encompasses both the voluntary and mandatory signature sessions while, in panel B, the voluntary and mandatory ethical treatments are separated. Consistent with Figure 4, Table 2 suggests a significant positive effect of observing an ethical statement on bids.

In Column 3 of panel A, the coefficient on $Ethical\ Statement$ is 0.345 and is statistically significant. This effect translates to bids that are on average nearly \$13 more. In Column 6 of panel A, buyers observing a signed ethical statement tend to bid nearly 38% higher. In the case of bidders observing the seller declining to sign, we observe a small positive coefficient of nearly 0.08 in Column 6 of both panels but the coefficient is not statistically significant at conventional levels. Turning to panel B, as shown in Columns 3 and 6, the positive effect of a seller signing an ethical statement on a buyer's bid seems to be stronger in the mandatory setting with a statistically significant coefficient of 0.452. In the voluntary setting, the coefficient of 0.327 and is significant at 10% level.

[8] In untabulated analysis we conduct two-by-two mean comparisons of bids between ethical treatments (mandatory vs voluntary) as well as between each ethical treatment and no treatment, using a t-test. We find : (1) statistically significantly higher average bid in the mandatory ethical treatment compared to the no ethical treatment setting, (2) no statistical significance in buyers average bid between mandatory and voluntary treatments or between voluntary and no treatment. However, these initial results can be driven by difference in amounts reported by the sellers. Therefore, we rely on the multivariate regressions results.

Table 2: Regression Analysis of Bids

Panel A	*Log Bid*					
	(1)	(2)	(3)	(4)	(5)	(6)
Ethics: Ethical Treatment	0.381**	0.376**	0.345*			
	(0.181)	(0.181)	(0.179)			
Ethics (=1) : Signed Ethical Statement				0.412**	0.411**	0.378**
				(0.181)	(0.182)	(0.181)
Ethics (=2) : Declined to Sign				0.125	0.093	0.083
				(0.221)	(0.222)	(0.214)
Observations	2,694	2,694	2,694	2,694	2,694	2,694
Adjusted R-squared	0.015	0.024	0.096	0.018	0.026	0.098
Mean. Bid	40.66	40.66	40.66	40.66	40.66	40.66
Controls	No	No	Yes	No	No	Yes
Round FE	No	Yes	Yes	No	Yes	Yes
Panel B	*Log Bid*					
	(1)	(2)	(3)	(4)	(5)	(6)
Mandatory : Signed Ethical Statement	0.467**	0.458**	0.452**			
	(0.211)	(0.213)	(0.215)			
Voluntary : Signed Ethical Statement				0.357*	0.362*	0.327*
				(0.196)	(0.197)	(0.191)
Voluntary : Declined to sign				0.125	0.099	0.088
				(0.221)	(0.222)	(0.213)
Observations	1,753	1,753	1,753	1,938	1,938	1,938
Adjusted R-squared	0.023	0.027	0.092	0.012	0.022	0.108
Mean. Bid	40.7	40.7	40.7	38.53	38.53	38.53
Controls	No	No	Yes	No	No	Yes
Round FE	No	Yes	Yes	No	Yes	Yes

Notes - This table summarizes the OLS regressions results of the amounts of bids by the buyers on : ethical treatment (whether voluntary or mandatory), and sellers choices regarding the signature of ethical statement in the voluntary treatment. The baseline is comprised of the bids in the sessions with no ethical treatment. In panel A the ethical treatment combines both the voluntary and mandatory signature sessions, while in panel B the voluntary and mandatory ethical treatments are seperated. The regressions are based on the sample of buyers only. *Log Bid* is the natural log of one plus the buyer's bid. *Ethics* is an indicator variable that takes the value of 1 if the seller is presented with an ethical statement to sign (voluntary or mandatory treatment); zero otherwise. *Ethics* is further decomposed into Signed *Ethical Statement* and *Declined to Sign*; indicator variables that take the values of 1 and 2 when the seller signs or declines, respectively; zero otherwise. Controls is a vector comprised of individual buyers' characteristics (age, gender, risk preference) as well as the reported amount. Round FE indicates whether the estimates are within rounds. The standard errors are clustered at the buyer level, and reported in parentheses below the coefficients. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

3.4 Trust and Ethics

We provide additional exploratory analysis on the relation between trust and ethics, consistent with Arrow's conjectured link between widely-accepted ethical codes and the trust required for efficient market transactions (Arrow 1973). To examine this relation, we utilize the trust questions obtained from our post experiment survey. We create a factor reflecting the aggregate trust of each participant and partition the sample of participants based on an aggregate trust score above or below the median trust score. We then proceed to re-estimate the models specified in Columns 3 and 6 from panel A of Table 1 for the sellers and panel A of Table 2 for the buyers on the partition samples based on trust.[9] Table 3 provides the results, with Columns 1-4 for sellers and 5-8 for buyers.

The effect of the ethics treatment on reports is not statistically different from zero across the two sub-samples based on trust level. However, in the case of the buyers, low-trust subjects choose higher bids in the ethical statement treatment. In contrast to high-trust buyers who exhibit low levels of skepticism regarding sellers' reports even absent ethical statement, the low-trust buyers raise their level of confidence regarding sellers' reports only if an ethical statement is signed. Specifically, the coefficients in the high-trust samples are 0.623 and 0.639 when considering the treatment based on a mandatory ethical statement and a voluntary ethical statement, respectively. These coefficients are statistically significant at 1% level, while the coefficients in the high-trust sub-sample are statistically insignificant and economically small.

4 Structural Estimation

4.1 The Seller's Problem

In this section, we first provide evidence from structural estimates that sellers do not behave as profit maximizers (i.e., $e \neq 1$ for most types), and there is a significant amount of uncertainty

[9] Due to limitation in sub-sample sizes following the partitioning based on trust level, we do not further separate the voluntary ethical statement signature treatment from the mandatory treatment as we did in our previous analyses

Table 3: Regression Analysis: Relation between Ethics and Trust

	Log Report (Sellers)				*Log Bid (Buyers)*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ethics: Ethical Treatment	0.041	0.020			0.623***	0.044		
	(0.052)	(0.050)			(0.218)	(0.282)		
Ethics (=1) : Signed Ethical Statement			0.052	0.008			0.639***	0.079
			(0.050)	(0.051)			(0.220)	(0.285)
Ethics (=2) : Declined to Sign			-0.040	0.108			0.459*	-0.191
			(0.153)	(0.080)			(0.260)	(0.330)
Observations	337	340	337	340	1,352	1,342	1,352	1,342
Adjusted R-squared	0.757	0.801	0.757	0.802	0.156	0.071	0.157	0.073
Mean Non Log Outcome	53.35	54.31	53.35	54.31	41.7	39.62	41.7	39.62
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Round FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Subsample	Low Trust	High Trust	Low Trust	High Trust	Low Trust	High Trust	Low Trust	High Trust

Notes - This table summarizes the OLS regressions results of the sellers' reports and buyers' bids on : ethical treatment (whether voluntary or mandatory) and sellers' choices regarding the signature of ethical statement in the voluntary treatment. Each of the samples of sellers and buyers is split between High- and Low-Trust subsamples based on the median level of trust measured in the survey. The baseline is comprised of the amounts reported or bids in the sessions with no ethical treatment. *Log Report* is the natural log of one plus the seller's report. *Log Bid* is the natural log of one plus the buyer's bid. *Ethics* is an indicator variable that takes the value of 1 if the seller is presented with an ethical statement to sign (voluntary or mandatory treatment); zero otherwise. *Ethics* is further decomposed into *Signed Ethical Statement* and *Declined to Sign*; indicator variables that take the values of 1 and 2 when the seller signs or declines, respectively; zero otherwise. Controls is a vector comprised of individual seller characteristics (age, gender, risk preference) as well as the true asset value. Round FE indicates whether the estimates are within rounds. The standard errors are clustered at the subject level, and reported in parentheses below the coefficients. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

across ethical preferences. We rely on equation (5) and first estimate the pricing function $\hat{P}_i(.)$ non-parametrically across four groups conditioning on the information available to buyers: treatments without ethical statement ($i = m_0$), with a mandatory ethical statement ($i = m_1$), with an ethical statement declined by the seller ($i = s_0$), and with an ethical statement signed by the seller ($i = s_1$).[10] We use a kernel instead of partitions (as in the reduced-form model) because the slope of the pricing function needs to be estimated, requiring more regularity than is possible with partitions.

[10]We use the optimal cross-validated bandwidth with the Epachnikov kernel, with a bandwidth of 18.3 for $i = ns$ ($R^2 = 0.59$), 307.7 for $i = s_0$ ($R^2 = 0.38$), and 16.2 for $i = s1$ ($R^2 = 0.57$). Confidence intervals were obtained by bootstrapping with 300 repetitions.

We also focus here on the winning bid because this is the payoff statistic relevant to the seller.

As noted earlier, in many first-price auction experiments (and ours is no exception), experimental subjects tend to overbid (Goeree et al. 2002; Crawford and Iriberri 2007), which could result from strategic errors or buyers deriving a supplementary utility from winning. For our research design below, we first take the sellers' perspective, which need not presume that buyers are following equilibrium play. For estimates relating buyers, we shall use buyer's estimate of the seller's valuation (which we induce in the experiment) to better pin down buyers' beliefs separately from how they bid. Nevertheless, as noted earlier, overbidding further tends to penalize the winning buyer and this effect is stronger with more heterogeneity about trust. We drop three observations with reports below zero, as they may indicate exploration and learning by certain sellers.[11]

Figure 5: Price response to reports, +/- one standard-error.

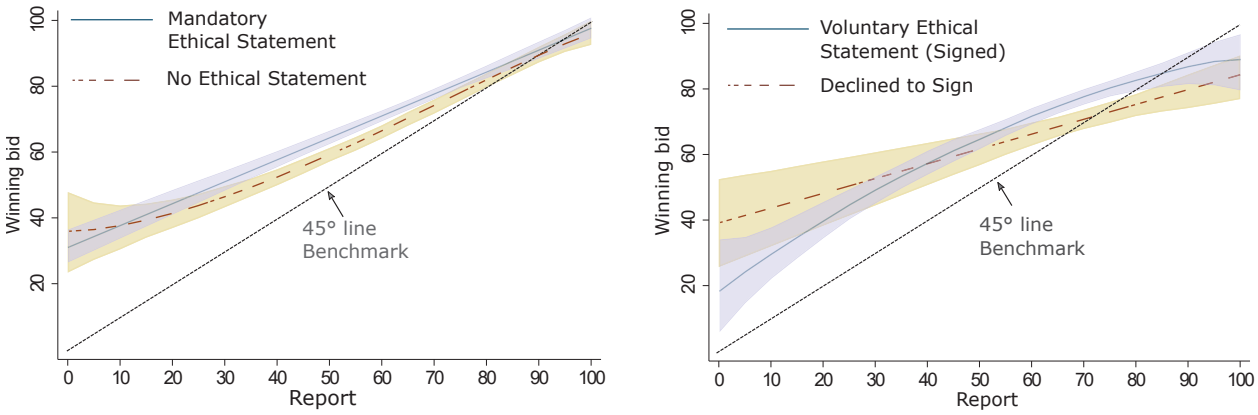

(a) Mandatory vs. No Ethical Statement

(b) Voluntary Signature vs Declined to Sign

Notes - This figure shows the estimated selling price conditional on asset value reported in the mandatory vs no ethical statement settings as presented in Panel (a); and in the voluntary setting when separating sellers with a signed ethical statement from those who declined to sign as presented in Panel (b). The 45-degree° line represents buyers bidding the reported value.

Figure 5 plots the estimated selling price over the mandatory (Panel (a)) and voluntary (Panel (b)) treatments, separating sellers with versus without a signed ethical statement. As in the reduced-form analysis, buyers become more skeptical over higher reports. In further tests, we find that

[11]This does not affect our results: our rationale is that reports below zero may be very different in nature from reports that are small but positive, thus, the price function may have to be estimated separately for these reports, rather as a part of a kernel estimate that smoothes small positive reports with negative reports.

signing an ethical statement reduces skepticism: prices are higher with a signed ethical statement for reports over \$25 with mandatory ethical statements and reports between \$65 and \$80 with a voluntary ethical treatment.

Figure 6: Ethical types by observation.

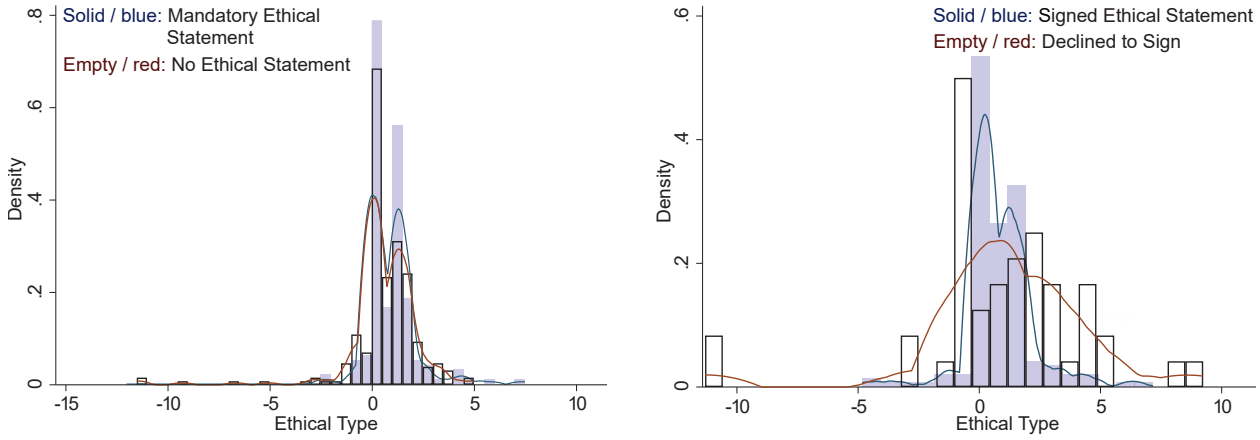

(a) Mandatory vs. No Ethical Statement

(b) Voluntary Signature vs Declined to Sign

Notes - This figure presents the distributions of inferred ethical coefficients over Mandatory vs. No Ethical Treatment in Panel (a); Voluntary Signature vs. Declined to Sign in Panel (b).

In Figure 6, we plot the histogram of inferred ethical coefficients over the mandatory and voluntary treatments, using the no ethical treatment setting and scenario where the seller declined to sign as respective benchmarks. Over all treatments, the mode of the distribution is located at zero, suggesting a modal preference for truthful reporting, in sharp contrast with the payoff-maximizing strategy $e = 1$. As we can see in the histogram, some sellers reported downwards (which one might attribute to an attempt to gain credibility).

The distribution of types conditional on no ethical statements or a signed ethical statements are similar, and features a bimodal distribution with a second mode located roughly at one, indicating a mass of observations consistent with maximizing economic payoffs. When declining to sign, we find an increase in misreporting, especially at very high or very low levels of misreporting. About 42% of sellers who declined to sign have an estimated $e > 2$, which indicates a sub-optimally aggressive reporting strategy (relative to monetary payoffs), versus only 16% and 9.9% in the no

ethical statement treatment and with a signed statement, respectively.[12] Interestingly, the distribution of ethical types in the models with mandatory signing and no ethical statements are very similar. One explanation for heterogeneity among observations may be that some of the buyers may have been experimenting to learn the pricing function. Subjects were given practice sessions and excluding the first five paid sessions does not alter these insights. Another explanation is that the process of strategic play is noisy, so that the wide distribution of e at the observation-level may reflect inherent difficulties in higher-level reasoning (Corgnet et al. 2018).

Figure 7: Ethical types averaged by subject

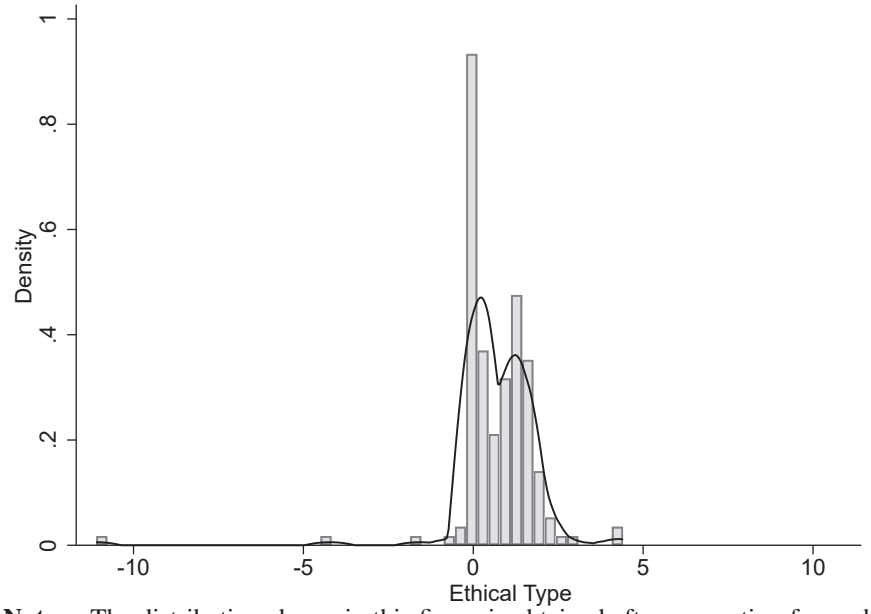

Notes - The distribution shown in this figure is obtained after computing for each subject the median inferred ethical type. The distribution takes into account all settings (mandatory/voluntary ethical statement and no ethical statement).

That many observations feature seemingly sub-optimal play $e \notin [0, 1]$ may also reflect noise. To average out this noise component, in Figure 7, we plot a histogram using the median e per subject. We continue to find strong evidence of large heterogeneity across sellers, with a bimodal distribution featuring slightly over a third (37%) of all subjects preferring truthful reporting (defined as $e \in (-0.25, 0.25)$) and 14% with behavior roughly consistent with maximizing payoffs

[12]This suggests that many subjects declining to sign may not have been playing the experiment seriously to maximize payoffs but may have been driven other considerations, such as deriving utility for surprising the buyer with extreme misreporting even if this is detrimental to them.

($e \in [.75, 1.25]$). The median (average) type is .69 (.67), which falls somewhere in-between truthful reporting and maximizing economic incentives. Approximately 63% (91%) have an estimated type $e \in [-0.25, 1.25]$ (in $[-0.25, 2]$), implying that some of the heterogeneity at the observation level with high levels of misreporting may be partly due to noise.

Table 4 documents differences between seller types, under three specifications, (1) ordinary least squares, (2) random effects, and (3) fixed effects.[13] The problem presents two difficulties. First, types are estimated from the optimality restriction in equation (5), which involves a noisy first-stage estimation of the pricing function; obtaining standard-errors thus requires consideration of first-stage standard errors. Second, Figure 6 suggests that some observations are outliers. To address these problems, we derive standard-errors by bootstrap, drawing a new subsample, re-estimating the price response, and then winsorize each bootstrap set of e estimates at the 5% level. We find across all treatments an average type of approximately $e \approx 0.6$ at mid-distance between truthful reporting and profit maximization. Consistent with the reduced-form analysis, we are unable to find a significant effect of signing an ethical statement reporting, relative to the baseline of no ethical statement. The parameter estimate on signing an ethical statement is positive but small, amounting to an effect on e of roughly 0.1, which would map to roughly $\$1$ of incremental bias in equation (3) and is much less than the between-subject variation.

A follow-up question is (a) whether sellers self-select into signing or not signing the ethical treatment as a function of their types, and (b) whether signing an ethical statement changes sellers' types by imposing a different psychological cost of misreporting similar to inducing moral thinking (Mazar et al. 2008). Naturally, these theories are not incompatible, and sellers may self-select while also considering the effect of signing on their preference as a commitment device. We address this question by considering the following implication of the pure self-selection theory: if signing an ethical statement does not directly affect preferences, the distribution of types across

[13] Subjects were assigned to a single treatment, so (3) can only be used to estimate the effect of signing the ethical statement in the voluntary treatment. We considered using subjects across treatments but one problem is that their observation of a prior treatments would affect their behavior in the next treatment, by making some choices more salient. We had no clear hypothesis about the consequences of this effect, so we did not run any experiment with within-subject treatment variation.

Table 4: Regression Analysis of Ethical Types

	Ehical type (Sellers)		
	(1) OLS	(2) Random effects	(3) Fixed effects
Mandatory : Signed Ethical Statement	0.111	0.068	
	(0.176)	(0.168)	
Voluntary : Signed Ethical Statement	0.071	0.021	
	(0.182)	-0.167	
Voluntary : Declined to sign	0.763*	0.701*	0.748*
	(0.427)	(0.407)	(0.443)
Trend (Round #)	0.013	0.013*	0.015
	(0.008)	(0.008)	(0.009)
Observations	690	690	690

Notes - This table shows differences in sellers' types (with respect to their levels of ethics) relative to the baseline of no ethical treatment, considering mandatory and voluntary treatments separately. The estimations are conducted using OLS (column 1), random effects (column 2), and fixed effects (column 3) at the seller level. The variable *Ethical Statement Signed* is an indicator that takes the value of one when in either the voluntary or mandatory settings the seller signs the ethical statement; zero otherwise. The variable *Declined to Sign* is an indicator that takes the value of one when in the voluntary setting the seller chooses not to sign the ethical statement; zero otherwise. *Trend* is a control variable equals to the round number to account for any learning effects. The standard errors are clustered at the seller level, and reported in parentheses below the coefficients. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

the three treatments, mandatory no statement, mandatory signed statement, and voluntary (pooling signers and non-signers) should be the same, because the presence of the ethical statement does not change the distribution of types. We provide various tests to compare the distributions under each treatment, and the resulting p-values, in Table 5.

When comparing treatments, we find very small differences between the distribution of types with a mandatory ethical statement versus no ethical statement, echoing prior observations that an ethical statement, if mandatory, does not seem to affect preferences. We observe a difference in the voluntary treatment, which features more mass over extreme types. In various hypotheses test assessing differences between distributions, we are unable to reject different distributions over a mandatory ethics statement and no statement (with both Kolmogorov-Smirnov and Wilcoxon

tests).[14] Across tests, we document some evidence that the distribution of types under the voluntary treatment is different from the treatment without an ethics statement. This suggests that, indeed, the *voluntary* statement changes ethics preferences and is not just driven by self-selection.

Table 5: Comparison of Type Distribution by Treatment

	KS, (1)>(2)	KS, (1)<(2)	KS, equal	Wilcoxon	Mean
Mandatory (1) vs. Voluntary (2)	0.249	0.011	0.023	0.497	0.728
Mandatory (1) vs. No Ethics (2)	0.113	0.218	0.226	0.182	0.030
Voluntary (1) vs. No Ethics (2)	0.038	0.998	0.077	0.064	0.028
Voluntary (1) vs. All other (2)	0.006	0.702	0.012	0.123	0.084
Voluntary Sign (1) vs. Decline (2)	0.703	0.000	0.000	0.005	0.025

Notes - The numbers in this table represent the p-values of relevant tests statistics: Kolmogorov-Smirnov test for a one sample greater than the other or equal distributions, Wilcoxon rank-sum test, and t-stat for equality of means. Populations refer to all treatments with mandatory ethics statement("Mandatory"), no ethics statements ("No Ethics"), voluntary ethics statement including both signers and non-signers ("Voluntary"), and, within the voluntary treatment, voluntary signers ("Voluntary Sign") and sellers declining to sign ("Decline").

4.2 The Buyers' Perspective

Thus far, we have focused on sellers' reporting choices, considering the *actual* distribution of ethical types and how they perceive buyers. In what follows, we examine in more detail buyers' *perception* of sellers' types. We begin with a simple introductory question: are buyers able to accurately infer the seller's information, as would be predicted by the rational model? To answer this question, we asked in the experiment buyers to provide their best assessment of the buyers' signal and compensate them as a function of a deviation from actuals: this data allows us to remove any benefit from winning the auction that could affect actual bids.

[14]The KS test statistic may here be inaccurate because there is are ties over types estimated at exactly $e = 0$. While, in principle, the p-value of the test statistic do not hold over discrete distributions or mixtures, it could be that the "true" distribution of types is continuous, but there are very small differences over zero that are rounded to $e = 0$.

Figure 8: Comparison of Prediction Errors.

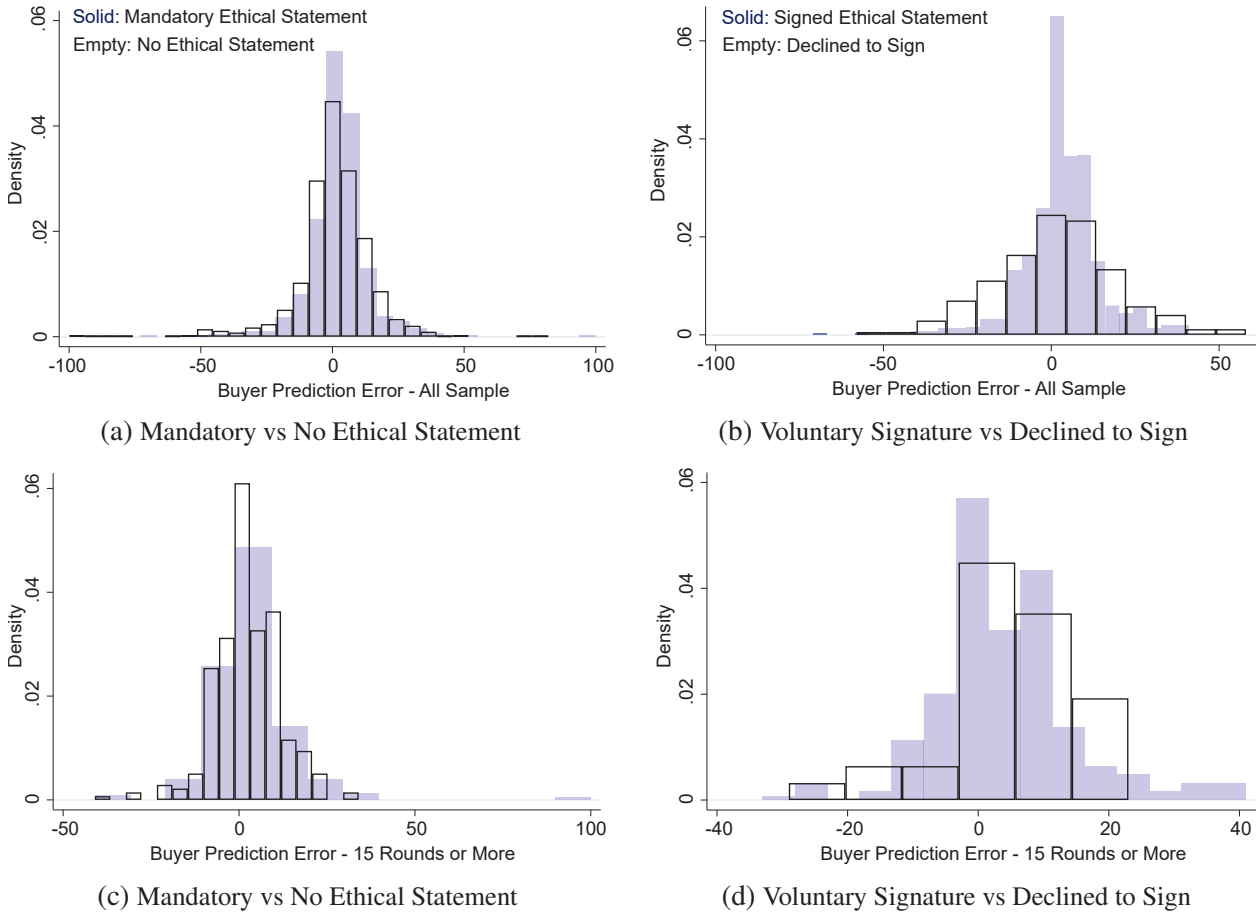

(a) Mandatory vs No Ethical Statement

(b) Voluntary Signature vs Declined to Sign

(c) Mandatory vs No Ethical Statement

(d) Voluntary Signature vs Declined to Sign

Notes - This figure shows the distribution of buyers' prediction errors accross different treatments and without treatment. Panels (c) and (d) complement panels (a) and (b) respectively by specifically focusing on experienced buyers (those having played 15 rounds or more) to elicited any learning effects.

First, we consider the buyers' prediction accuracy relative to actual bias behavior in the experiment, defined as the difference between the prediction and the true asset value. Note that buyers are incentivized by incurring a monetary cost proportional to the absolute difference between prediction and value regardless of whether they win the auction. In Figure 8, we plot the buyers' prediction errors over each treatment. On average, relative (absolute) prediction errors are \$2.3 (\$7.45), \$.5 (\$9.6), \$3.0 (\$8.1), and \$1.50 (\$14.0) in the mandatory ethics, no ethics, voluntary signing and voluntary decline to sign samples, respectively. Overall, we find evidence of a "strategic" wisdom of the crowd effect such that, *on average*, the combined estimate of all buyers is able to reasonably correct the bias. Interestingly, because players in the experiment have no pri-

vate information about fundamentals, this effect is not driven by an aggregation of fundamental information but how the crowd appears to average out strategic errors. Practically, concerns about misreporting and the creation of regulatory bodies such as the SEC are responses to suspected behavioral errors by investors (Choi and Pritchard 2003; Bourveau et al. 2020); hence, these findings suggests that these functions may have a more important role in thin markets (Gerakos et al. 2013).

The average absolute prediction errors and the distribution of errors in Figure 8 strongly reject that subjects are individually able to infer the true value. Errors are much larger when sellers decline to sign the ethical statement, and to a lesser extent, absent any ethical statement. We further conduct two-by-two mean difference tests. In the voluntary treatment, average absolute errors are significantly higher, at the 1% level, than with a mandatory ethical statement. However, the economic magnitude of the difference is small, at \$1.8. We find no significant difference between errors in the voluntary treatment versus without ethical statement. Our conclusion is two-fold: subjects do not appear to accurately infer the seller's information but, except in the case where a seller refuses to sign, the effect of ethical statements on buyers' inference is small.

One explanation for this failure of rational inference may be that buyers experimented beyond the training sessions (which were not included in the above analysis). To test this, we repeat the analysis using only the last 15 rounds. We find that results in the experienced subsample are very similar, both in terms of distribution and average errors. On average, relative (absolute) prediction errors are \$3.0 (\$7.3), \$2.3 (\$7.1), \$4.2 (\$8.1), and \$3.8 (\$8.3) in the mandatory ethics, no ethics, voluntary signing and voluntary decline to sign samples, respectively. The main difference is that absolute errors when subjects refuse to sign are greatly reduced, which we attribute to subjects learning skepticism and is consistent with similar results in disclosure applications (Dickhaut et al. 2003; Jin et al. 2021; Bourveau et al. 2020).

Figure 9: Ethical Types - Buyers' perception.

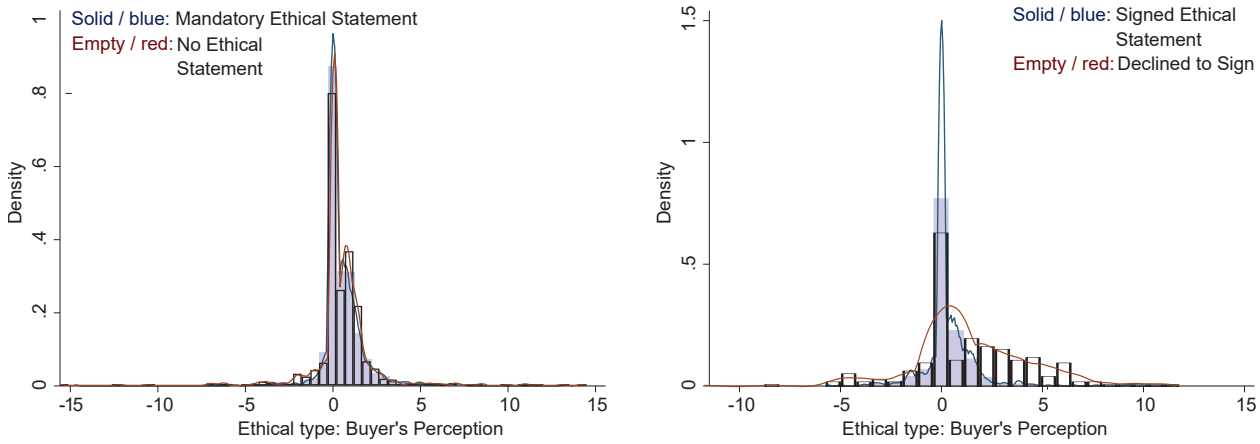

(a) Mandatory vs No Ethical Statement

(b) Voluntary Signature vs Declined to Sign

Notes - This figure shows the distribution of the estimated perception of buyers regarding sellers' types in both the mandatory ethical statement and no ethical statement settings (panel (a)) as well as in the voluntary ethical statement treatment differentiating scenarios where the statement is signed from those where sellers decline to sign (panel (b)).

Our second approach is to compare the actual distribution of seller type to that of inferred seller type by the buyers. We exploit equation (5) using the bias $r - \hat{v}$ reported by the buyer, i.e., where $\hat{v}$ is the prediction. In Figure 9, the distribution of inferred ethics is similar to that of the actual ethics. We still see a mass point of honest sellers but the peak around profit-maximization $e = 1$ is less clear-cut, perhaps because it is an indirect process for a buyer to infer ethics than for a seller to play according to their preference. Note that this result is also counter-intuitive: buyers' inference of ethical types should be *less* dispersed than actual types since their beliefs are noisy relative to actual types. Instead, we observe that buyers disagree on sellers' ethics and the outcome of their disagreement is reflected in Figure 9. The result further emphasizes that ethical noise is not solely caused by actual ethics dispersion, but also by disagreeing perceptions about ethics, with limited common-knowledge of ethics differences and such that neither sellers or buyers have accurate higher-order beliefs about what other parties believe about ethics. Nevertheless, consistent with results on prediction errors, the average perceived type is very similar to the actual type as inferred in Figure 6.

A possible behavioral explanation for this result is that subjects project their own type as sell-

ers when they act as buyers, noting prior evidence showing that individuals project their social preferences on others (Ambuehl et al. 2019). To test this hypothesis, we examine the difference per subject between their median inferred type as sellers and their median perceived type over the seller when they act as buyers. However, while the correlation is positive (.08), it is small and insignificant at conventional levels. In summary, we do not find evidence that subjects treat others in the manner they expect to be treated.

4.3 Rational Expectations with Noisy Ethics

From the seller's and buyer's perspectives, we examined experimental results assuming that neither buyers nor sellers necessarily maximize their (ethical) preferences if they had informed common-knowledge of the distribution of ethical types. In this section, we explore the potential benefits, or side effects, of greater education about ethical noise that could facilitate convergence of buyers and sellers toward correct beliefs. We ask how the right level of skepticism calibrated to actual preferences may help reduce noise and improve communication. Note that an answer to this question is a-priori unclear, as greater rationality will make the buyer's belief more accurate, holding the seller's actions constant, but the resulting change in price will ambiguously affect the seller's strategy.

To analyze this counter-factual, we solve for the rational expectations equilibrium of the model, taking the estimated preferences of the seller and assuming that both buyers and sellers choose mutual best responses.[15] Because the process for true values and ethics is neither bounded nor normally-distributed, the problem does not have a linear equilibrium as in Fischer and Verrecchia (2000), so we develop below a numerical iterative procedure to solve the model. We construct a uniform grid of values (v, e) by stacking 100 times a grid $v \in \{0, \ldots, 100\}$ next to 100 random permutations of the grid $e \in E = \{e_1, \ldots, e_J\}$, for a total of $101 \times J$ sample points (v_i, e_i). Starting from $P_0(r) \equiv r$, we then solve for a best response report $R_0(v_i, e_i)$ using P_0 in equation(1) and finding the global maximum over reports $r \in \{0, 1, \ldots, 120\}$.[16] We then update the price function

[15] We use estimated the median e by subject, as shown in Figure 7, to reduce the noise in the estimation.

[16] To compare with the fully-separating equilibrium in (4), we allow for reports greater than 100; however, the

using a learning rate ρ as

$$P_1(r) \equiv \rho \mathbb{E}(\tilde{v}|R_k = r) + (1-\rho)P_0(r), \tag{8}$$

where $\mathbb{E}(\tilde{v}|R_k = r)$ is obtained by averaging over the realized R_k. We repeat this procedure replacing P by P_k in equations (1) and (8), and updating P_{k+1} until the average difference $|P_k - \mathbb{E}(\tilde{v}|R_k = r)|$ is small.

It may be intuitive to use $\rho = 1$ in (8), which corresponds to iterated best responses as the pricing function is repeated updated to the seller's strategy while the seller's report is re-optimized to the new pricing function. However, numerically, this algorithm does not converge. For two given reports close to each other r_1 and r_2 that are not yet at the equilibrium pricing function, $P_k(r_2)$ may be relatively high relative to $P_k(r_1)$, so that many unethical sellers prefer r_2. At the next iteration, the report r_2 will be heavily discounted, making $P_{k+1}(r_1)$ relatively high relative to $P_{k+1}(r_2)$. We find that the process oscillates and fails to converge. Decreasing the learning rate forces the algorithm to continuously change the pricing function and prevents this. We found that $\rho = 0.01$ ensured convergence, at which point the conditions are verified at any desired precision.

In Figure 10, we plot the equilibrium reporting function. We show that this function features some degree of skepticism typically midway between truthful reporting and payoff-maximization, except for very low reports below 10 which tend to be anticipated to be under-reporting. As reports become more aggressive, the pricing function tends to become similar to common-knowledge of no ethics, which shows how strategies update that a higher report is more likely to be from a less ethical preference.

analysis is very similar if we bound reports by 100 since only values very close to 100 will attain this upper bound.

Figure 10: Equilibrium Pricing Function under Random Ethics (Actual data)

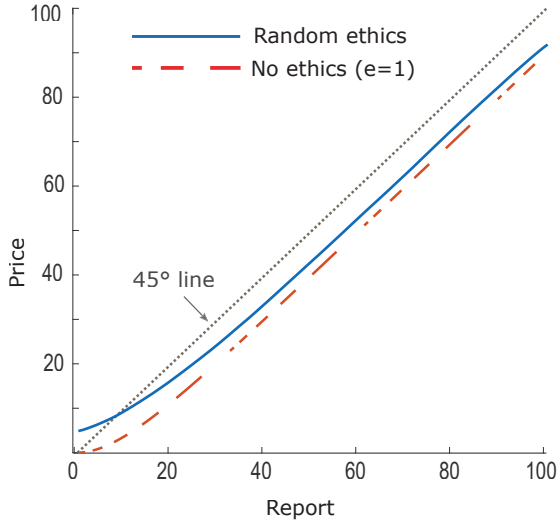

Notes - This figure compares the actual price function considering the reported asset value as input (curve referenced "Random ethics", generated from the results of the experiment) to the counter-factual price function (curve referenced "No ethics (e=1)" obtained by simulation from rational expectations model). The 45-degree° line represents truthful reporting.

In Table 6, we compare pricing errors, bias, and seller welfare in the experiment, versus the rational expectations solution (the buyer's trading loss can be obtained by subtracting the average pricing error).

Table 6: Counter-factual Analyses

	avg. price error	avg. | price error |	avg. bias	seller util (e=1)	avg. seller util (all e)
Experiment (all treatments)	1.79	8.90	4.94	57.41	42.58
Nash ethics noise	0.26	6.30	6.22	48.57	36.45
Nash ethics noise (cost x 2)	0.29	3.46	3.21	49.47	36.11
Nash all maximizers	0.03	0.09	9.48	45.44	45.44

Notes - This table summarizes the average price errors (both relative and absolute value), average seller's bias, and average seller's utility (considering all types of sellers on the one hand and sellers of type e=1 on the other hand). price error is the difference between buyer's estimation and true asset value. seller's bias is the difference between seller's report and asset value. Seller utility takes into account both seller's monetary gain and intrinsic misreporting cost. The first row presents the averages obtained from the results of the experiment. In the second and third rows, the averages are obtained from the results of simulations, and serve as counterfactuals when using a rational expectation model that takes into account ethical noise. The difference between the second and third rows is that, in the third row, misreporting cost is multiplied by 2 in the model used to generate the results in the third row. The fourth row provides the averages based on simulation when considering a model in which it is common-knowledge that all sellers are payoff-maximizing as counterfactual.

Subjects in the experiment feature higher absolute pricing errors, i.e., lower price efficiency, than what they would have achieved under rational expectations even though, in equilibrium, buyers do not observe ethical preferences e and cannot fully extract the asset values. This effect does not occur because of an average bias: by construction, the price is on average equal to the true value in the Nash equilibrium (Column 1) but it remains very small even in the experiment ($1.79 average bias). However, the absolute pricing error is reduced by about 30% in the equilibrium (Column 2), showing that learning the optimal skepticism makes price more informative by correcting extreme over and under pricing. This occurs even though the average bias increases by 26% from $4.94 to $6.22 (Column 3), as sellers adapt to greater skepticism by biasing more. In summary, more rationality jointly increases predictable bias while reducing noise in inferring true asset values. Interestingly, the seller's utility of both a payoff-maximizing seller and an average seller decrease (Columns 4 and 5), which confirms that rational pricing reduces opportunities to exploit buyers' excess trust.

We also examine the effect of greater ethical noise and compare these results to the model in which it is common-knowledge that all sellers are payoff-maximizing ("all maximizers"). As expected, we find that a greater misreporting cost (multiplied by 2) reduces predictable bias and pricing errors, roughly halving them. When preferences are uncertain, Fischer and Verrecchia (2000) prove that greater reporting flexibility may, under certain conditions, benefit sellers even though buyers price-protect, because sellers can better adjust their reporting strategy as a function of e. We assess this conjecture in Columns 4 and 5 and find that sellers are (marginally) better-off with higher misreporting costs, which is also true in the model without uncertainty about ethics.

5 Conclusion

This study examines the relation between ethical declarations and market transactions. Our results challenge the conventional view that enhanced ethical reporting inherently benefits markets. Along several treatments, ethics considerations are prevalent among participants in the experiment regardless of whether or not they sign an ethical statement, highlighting the significant role of

psychological costs of lying. However, contrary to popular beliefs, standard ethical interventions —had minimal impact on altering seller behavior beyond their baseline ethical conduct. Ethical statements only slightly modified reports except in cases where subjects opted out of signing. Interestingly, these ethical statement tended to be viewed as more credible by certain trusting buyers, which led them to overpay in response to sellers' reports. Unlike predictable biases, this phenomenon, termed "ethical noise," along with the difficulty buyers face in forming accurate expectations about sellers ethics, can reduce price efficiency compared to settings where sellers solely focus on maximize payoffs, which in turn ultimately disadvantages buyers.

While our experiment is unique in examining ethical declarations on market behavior, it is important to note several inherent limitations. Our inferences are derived from the revealed preferences of participants, observing how subjects act as sellers (via biases) and buyers (via how their estimate and prices). Subjects in the experiment are compensated based on how well they do monetarily, which allows us to disentangle their ethical preferences. While effective at extracting preferences, this experimental design is confined to the laboratory setting and might not capture all considerations in real-world ethical decision-making.

Bibliography

Abdel-Rahim, Heba Y, and Douglas E Stevens (2018) 'Information system precision and honesty in managerial reporting: A re-examination of information asymmetry effects.' *Accounting, Organizations and Society* 64, 31–43

Acemoglu, Daron, Ali Cheema, Asim I Khwaja, and James A Robinson (2020) 'Trust in state and nonstate actors: Evidence from dispute resolution in pakistan.' *Journal of Political Economy* 128(8), 3090–3147

Aghamolla, Cyrus, and Kevin Smith (2022) 'Strategic complexity in disclosure.' *Available at SSRN 3936805*

Ambuehl, Sandro, B Douglas Bernheim, and Axel Ockenfels (2019) 'Projective paternalism.' Technical Report, National Bureau of Economic Research

Arrow, Kenneth (2006) 'The economy of trust.' *Religion and Liberty* 16(3), 3

Arrow, Kenneth J (1973) 'Social responsibility and economic efficiency.' *Public Policy* 21(3), 303–317

Azrieli, Yaron, Christopher P Chambers, and Paul J Healy (2018) 'Incentives in experiments: A theoretical analysis.' *Journal of Political Economy* 126(4), 1472–1503

Barrios, John M, Efraim Benmelech, Yael V Hochberg, Paola Sapienza, and Luigi Zingales (2021a) 'Civic capital and social distancing during the covid-19 pandemic?' *Journal of public economics* 193, 104310

Barrios, John M., Efraim Benmelech, Yael V. Hochberg, Paola Sapienza, and Luigi Zingales (2021b) 'Civic capital and social distancing during the covid-19 pandemic.' *Journal of Public Economics* 193, 104310

Berg, Joyce, John Dickhaut, and Kevin McCabe (1995) 'Trust, reciprocity, and social history.' *Games and economic behavior* 10(1), 122–142

Bertomeu, Jeremy, Edward Li, Edwige Cheynel, and Ying Liang (2019) 'How uncertain is the market about managers' reporting objectives? evidence from structural estimation'

Bertomeu, Jeremy, Masako Darrough, and Wenjie Jason Xue (2013) 'Optimal conservatism with earnings manipulation.' *Conference of Contemporary Accounting Research, Fall 2013*

Beyer, Anne, Daniel A. Cohen, Thomas Z. Lys, and Beverly R. Walther (2010) 'The financial reporting environment: Review of the recent literature.' *Journal of Accounting and Economics* 50(2-3), 296–343

Bloomfield, Robert, and Maureen O'Hara (1999) 'Market transparency: who wins and who loses?' *The Review of Financial Studies* 12(1), 5–35

Bloomfield, Robert, Maureen Ohara, and Gideon Saar (2009) 'How noise trading affects markets: An experimental analysis.' *The Review of Financial Studies* 22(6), 2275–2302

Bourveau, Thomas, Matthias Breuer, and Robert C Stoumbos (2020) 'Learning to disclose: Disclosure dynamics in the 1890s streetcar industry.' *Available at SSRN 3757679*

Caskey, Judson, and Volker Laux (2013) 'Effects of corporate governance and managerial optimism on accounting conservatism and manipulation'

Charness, Gary, Uri Gneezy, and Brianna Halladay (2016) 'Experimental methods: Pay one or pay all.' *Journal of Economic Behavior & Organization* 131, 141–150

Choi, Stephen J, and Adam C Pritchard (2003) 'Behavioral economics and the sec.' *Stanford Law Review* pp. 1–73

Corgnet, Brice, Mark Desantis, and David Porter (2018) 'What makes a good trader? on the role of intuition and reflection on trader performance.' *The Journal of Finance* 73(3), 1113–1137

Crawford, Vincent P, and Nagore Iriberri (2007) 'Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions?' *Econometrica* 75(6), 1721–1770

Crockett, Sean, Vernon L Smith, and Bart J Wilson (2009) 'Exchange and specialisation as a discovery process.' *The Economic Journal* 119(539), 1162–1188

Da, Zhi, and Xing Huang (2020) 'Harnessing the wisdom of crowds.' *Management Science* 66(5), 1847–1867

Davidson, Bruce I, and Douglas E Stevens (2013) 'Can a code of ethics improve manager behavior and investor confidence? an experimental study.' *The Accounting Review* 88(1), 51–74

Dickhaut, John, Margaret Ledyard, Arijit Mukherji, and Haresh Sapra (2003) 'Information management and valuation: an experimental investigation.' *Games and Economic Behavior* 44(1), 26–53

Dye, Ronald A. (1988) 'Earnings management in an overlapping generations model.' *Journal of Accounting Research* 26(2), 195–235

Einhorn, Eti, and Amir Ziv (2008) 'Intertemporal dynamics of corporate voluntary disclosures.' *Journal of Accounting Research* 46(3), 567–589

_ (2012) 'Biased voluntary disclosure.' *Review of Accounting Studies* 17(2), 420–442

Elfenbein, Daniel W, Raymond Fisman, and Brian McManus (2019) 'Does cheap talk affect market outcomes? evidence from ebay.' *American Economic Journal: Applied Economics* 11(4), 305–326

Evans, John H. III, R. Lynn Hannan, Ranjani Krishnan, and Donald V. Moser (2001) 'Honesty in managerial reporting.' *The Accounting Review* 76(4), 537–559

Falk, Armin, and Nora Szech (2013) 'Morals and markets.' *science* 340(6133), 707–711

Fischer, Paul E., and Robert E. Verrecchia (2000) 'Reporting bias.' *The Accounting Review* 75(2), 229–245

Forsythe, Robert, Russell Lundholm, and Thomas Rietz (1999) 'Cheap talk, fraud, and adverse selection in financial markets: Some experimental evidence.' *The Review of Financial Studies* 12(3), 481–518

Frankel, Alex, and Navin Kartik (2019) 'Muddled information.' *Journal of Political Economy* 127(4), 1739–1776

Gerakos, Joseph, Mark Lang, and Mark Maffett (2013) 'Post-listing performance and private sector regulation: The experience of london's alternative investment market.' *Journal of Accounting and Economics* 56(2-3), 189–215

Glaeser, Edward L, David I Laibson, Jose A Scheinkman, and Christine L Soutter (2000) 'Measuring trust.' *The quarterly journal of economics* 115(3), 811–846

Gneezy, Uri, Agne Kajackaite, and Joel Sobel (2018) 'Lying aversion and the size of the lie.' *American Economic Review* 108(2), 419–453

Gneezy, Uri, Bettina Rockenbach, and Marta Serra-Garcia (2013) 'Measuring lying aversion.' *Journal of Economic Behavior & Organization* 93, 293–300

Goeree, Jacob K, Charles A Holt, and Thomas R Palfrey (2002) 'Quantal response equilibrium and overbidding in private-value auctions.' *Journal of Economic Theory* 104(1), 247–272

Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2008a) 'Trusting the stock market.' *the Journal of Finance* 63(6), 2557–2600

_ (2008b) 'Trusting the stock market.' *the Journal of Finance* 63(6), 2557–2600

Guttman, Ilan, Ohad Kadan, and Eugene Kandel (2006) 'A rational expectations theory of kinks in financial reporting.' *The Accounting Review* 81(4), 811–848

_ (2010) 'Dividend stickiness and strategic pooling.' *Review of Financial Studies* 23(12), 4455–4495

Hurkens, Sjaak, and Navin Kartik (2009) 'Would i lie to you? on social preferences and lying aversion.' *Experimental Economics* 12, 180–192

Jiang, Xu, Baohua Xin, and Yan Xiong (2022) 'The value of mandatory certification: A real effects perspective.' *Journal of Accounting Research*

Jin, Ginger Zhe, and Andrew Kato (2006) 'Price, quality, and reputation: Evidence from an online field experiment.' *The RAND Journal of Economics* 37(4), 983–1005

Jin, Ginger Zhe, Michael Luca, and Daniel Martin (2021) 'Is no news (perceived as) bad news? an experimental investigation of information disclosure.' *American Economic Journal: Microeconomics* 13(2), 141–173

_ (2022) 'Complex disclosure.' *Management Science* 68(5), 3236–3261

Kartik, Naveen, Marco Ottaviani, and Francesco Squintani (2007) 'Credulity, lies, and costly talk.' *Journal of Economic Theory* 134(1), 93–116

Kartik, Navin (2009) 'Strategic communication with lying costs.' *The Review of Economic Studies* 76(4), 1359–1395

King, Ronald R (1996) 'Reputation formation for reliable reporting: An experimental investigation.' *Accounting Review* pp. 375–396

Kreps, David M, Paul Milgrom, John Roberts, and Robert Wilson (1982) 'Rational cooperation in the finitely repeated prisoners' dilemma.' *Journal of Economic theory* 27(2), 245–252

Kristal, Ariella S, Ashley V Whillans, Max H Bazerman, Francesca Gino, Lisa L Shu, Nina Mazar, and Dan Ariely (2020) 'Signing at the beginning versus at the end does not decrease dishonesty.' *Proceedings of the National Academy of Sciences* 117(13), 7103–7107

Marco, Catherine A, and Gregory L Larkin (2000) 'Research ethics: ethical issues of data reporting and the quest for authenticity.' *Academic Emergency Medicine* 7(6), 691–694

Marinovic, Ivan (2013) 'Internal control system, earnings quality, and the dynamics of financial reporting.' *The RAND Journal of Economics* 44(1), 145–167

Mazar, Nina, On Amir, and Dan Ariely (2008) 'The dishonesty of honest people: A theory of self-concept maintenance.' *Journal of marketing research* 45(6), 633–644

Milgrom, Paul R. (1981) 'Good news and bad news: Representation theorems and applications.' *Bell Journal of Economics* 12(2), 380–391

Milgrom, Paul R, Douglass C North, and Barry R Weingast* (1990) 'The role of institutions in the revival of trade: The law merchant, private judges, and the champagne fairs.' *Economics & Politics* 2(1), 1–23

Mittendorf, Brian (2006) 'Capital budgeting when managers value both honesty and perquisites.' *Journal of Management Accounting Research* 18(1), 77–95

O'Brien, John, and Sanjay Srivastava (1991) 'Dynamic stock markets with multiple assets: An experimental analysis.' *The Journal of Finance* 46(5), 1811–1838

Sapienza, Paola, and Luigi Zingales (2012) 'A trust crisis.' *International Review of Finance* 12(2), 123–131

Sapienza, Paola, Anna Toldra-Simats, and Luigi Zingales (2013) 'Understanding trust.' *The Economic Journal* 123(573), 1313–1332

Sheremeta, Roman M, and Timothy W Shields (2013) 'Do liars believe? beliefs and other-regarding preferences in sender-receiver games.' *Journal of Economic Behavior & Organization* 94, 268–277

Smith, Adam (1776) '1976. an inquiry into the nature and causes of the wealth of nations.' *The Glasgow edition of the works and correspondence of Adam Smith*

Stein, Jeremy C. (1989) 'Efficient capital markets, inefficient firms: A model of myopic corporate behavior.' *Quarterly Journal of Economics* 104(4), 655–669

Stevens, Douglas E (2023) *In Search of a Moral Foundation for Capitalism: From Adam Smith to Amartya Sen* (Cambridge University Press)

van Dolder, Dennie, and Martijn J van den Assem (2018) 'The wisdom of the inner crowd in three large natural experiments.' *Nature Human Behaviour* 2(1), 21–26

Yang, Kaifeng (2009) 'Examining perceived honest performance reporting by public organizations: Bureaucratic politics and organizational practice.' *Journal of Public Administration Research and Theory* 19(1), 81–105

Online Appendix: Experiment Instructions

Program link: `https://stag-σklapp.herokuapp.com/`

Window A.0 - Instructions Quiz

Instructions Quiz

Before proceeding to the Practice Rounds, please answer the following questions regarding the instructions you just read:

1. In each round you'll be grouped with 4 other participants. Is it likely that you'll be grouped with the exact same participants in more than 1 round?

- ○ Extremely likely, I'll be grouped with the exact same participants in most rounds
--> That is incorrect, please try again
- ○ Fairly likely, I'll be grouped with the exact same participants every other round
- ◉ Extremely unlikely, I'll probably be grouped with the exact same participants only once
--> Correct! Chances of being grouped with the exact same group are extremely low

2. Are instructions going to be changing throughout the experiment?

- ◉ No, instructions stay the same throughout the experiment
--> Correct! Instructions stay the same and are repeated in most decision screens
- ○ Yes, new instructions will be provided every other round
- ○ Yes, instructions change in every round

3. The Asset Seller in each group is required to send a report (truthful or otherwise) to all potential Asset Buyers. What is this report about? :

- ○ It's about the color of the asset being auctioned
- ○ It's about the underlying value of the asset
- ◉ It's about the minimum price the Seller would accept
--> That is incorrect, please try again

4. In each round the value of the asset is known to the group Seller but not the potential Buyers. Something known to everyone about the asset value is:

- ○ That its value is always between L$25 and L$75
- ◉ That its value is always between L$0 and L$100
--> Correct! The asset value is unknown to potential Buyers, but always lies between L$0 and L$100
- ○ That its value is always between L$75 and L$150

5. Which of the potential Asset Buyers receives the value of the asset?

- ○ Whoever places the second highest bid
--> That is incorrect, please try again
- ○ Whoever places the bid closest to the underlying value of the asset
- ◉ Whoever places the highest bid
--> Correct! The highest bidder receives the value of the asset

6. The group's Asset Seller can always choose to issue an untruthful report. If so, what happens:

- ○ The Seller faces an increasing cost when misreporting
- ○ The Seller may get caught and expelled from the experiment
- ○ Nothing happens, there's no penalty for misreporting

Window A.1 - Seller's Screen:

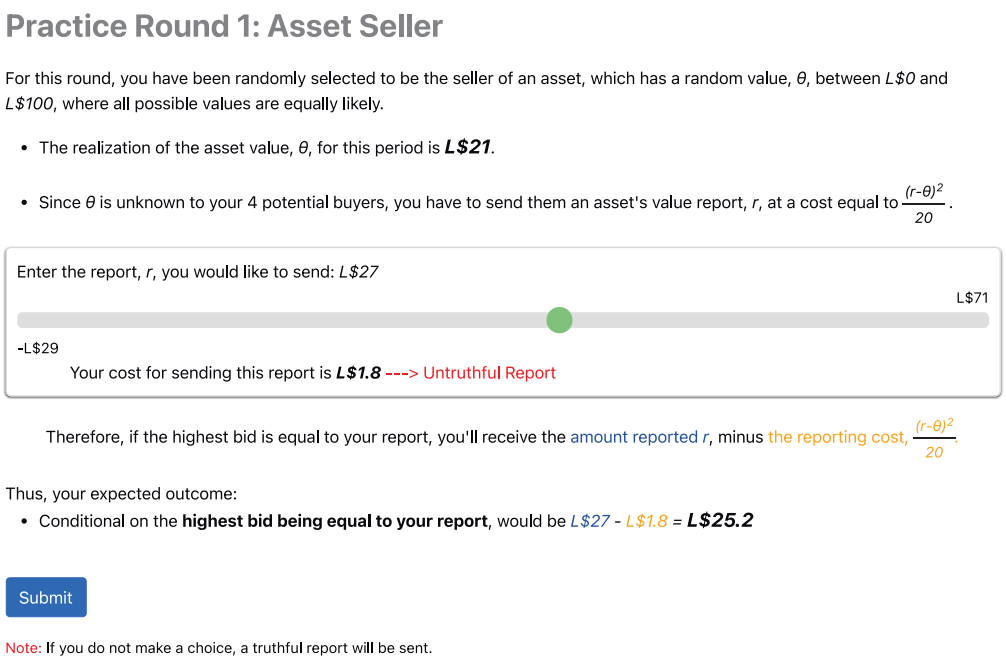

Note:
1. The subject can move the slider around to see what his reporting cost will be, plus the expected outcome conditional on the highest bid being equal to their report.
2. If the subject does not make a choice, the report defaults to the asset value.

Window A.2 - Ethical Treatment pop-up window:

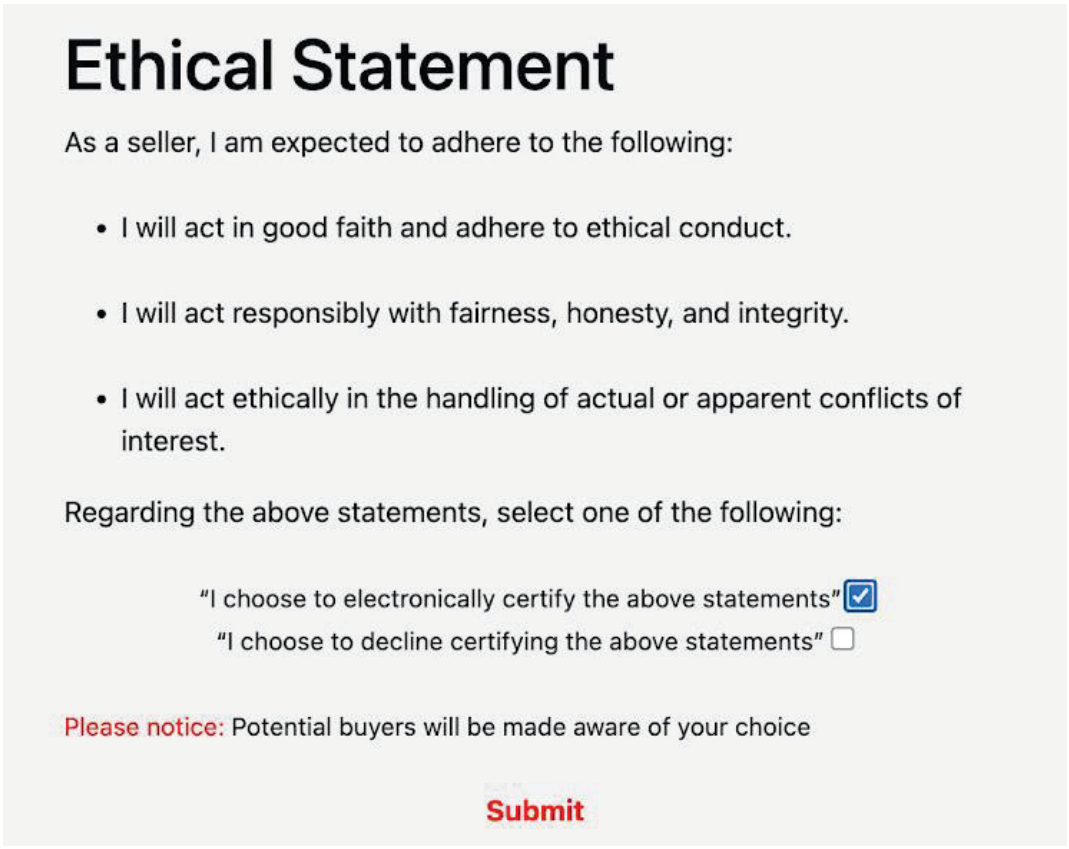

Window A.3 - Buyer's Screen:

Practice Round 1: Asset Buyer

For this practice round, you have been randomly selected to be 1 of 4 potential buyers of an asset, which has a random value, θ, between *L$0* and *L$100*.

- θ is unknown to you (and the other buyers), but the seller of the asset chose to report a value, r, of ***L$27***. This report may be inaccurate, but inaccuracy is costly for the seller (*see Cost Examples*)
- And the seller **chose to certify** the Ethical Statement (*see Ethical Statement*)
- You are endowed with *L$100* to use completely or partially to bid for the asset in this round.
- You'll have to make a prediction of the asset value, θ_p, which helps inform your decision and carries a prediction error cost of $|\theta-\theta_p|$.

Enter your prediction for the underlaying value of the asset, θ_p: ***L$ 38***

L$100

L$0

How much would you like to bid for this asset, b: ***L$ 27***

L$100

L$0

Therefore, if you turn out to be the highest bidder, you'll have your endowment minus your bid, *L$100-b* left, receive the value of the asset, θ, and pay your prediction error cost, $|\theta-\theta_p|$.

Thus, your expected outcome:

- Conditional on the **seller's report being correct**, would be *L$ 73+ L$27-L$ 11=**L$ 89***
- Conditional on **your prediction being correct**, would be *L$73+ L$38-L$0=**L$ 111***

Submit

Note: If you do not make a choice, your bid will default to zero and your prediction to the report.
Note: If there is a tie among the highest bidders, who takes the asset and pays for it will be randomly selected.

Window A.4 - Buyer's Result Screen:

Buyer's Results for Practice Round 1

For this practice round you were a potential buyer:

- The asset value, θ, was ***L$21***.
- Your asset value prediction, θ_p, was *L$28* so your prediction error cost is ***L$7***.
- Your bid of ***L$27* was the highest**, so you keep ***L$73*** out of your round endowment **and receive the asset value of *L$21***.

Thus, your payoff for this round is equal to ***L$21-L$7+L$73=L$87***

Next

Results including this round

Round	Your Role	Your Bid, b	Highest Bid	Your Asset Value Prediction, θ_p	Seller's Report, r	Ethical Statement	True Asset Value, θ	Your Round Payoff
1	Buyer	L$27	L$27	L$28	L$27	Certified	L$21	**L$87.0**

Window A.5 - Exit Questionnaire:

Exit Questionnaire

7. For each statement, select the option that best describes how much you agree or disagree with each statement.

Statement	How much you agree or disagree
"One should be very cautious with strangers"	○ Disagree strongly ○ Disagree ○ Neither agree or disagree ○ Agree ○ Agree strongly
"Most experts tell the truth about the limits of their knowledge"	○ Disagree strongly ○ Disagree ○ Neither agree or disagree ○ Agree ○ Agree strongly
"Most people can be counted on to do what they say they will do"	○ Disagree strongly ○ Disagree ○ Neither agree or disagree ○ Agree ○ Agree strongly
"These days, you must be alert or someone is likely to take advantage of you"	○ Disagree strongly ○ Disagree ○ Neither agree or disagree ○ Agree ○ Agree strongly
"Most salespeople are honest in describing their products"	○ Disagree strongly ○ Disagree ○ Neither agree or disagree ○ Agree ○ Agree strongly
"Most repair people will not overcharge people who are ignorant of their specialty"	○ Disagree strongly ○ Disagree ○ Neither agree or disagree ○ Agree ○ Agree strongly
"Most people answer public opinion polls honestly"	○ Disagree strongly ○ Disagree ○ Neither agree or disagree ○ Agree ○ Agree strongly
"Most adults are competent at their jobs"	○ Disagree strongly ○ Disagree ○ Neither agree or disagree ○ Agree ○ Agree strongly

Next