Ethics and Illusions: How Ethical Declarations Shape Market Behavior

Abstract

We study how ethical declarations shape market behavior using a laboratory setting where privately informed sellers issue reports before a first-price auction. Signing an ethical statement does not meaningfully reduce sellers' misreporting. Instead, it shifts beliefs: bidders place more weight on reports and bid more aggressively, moving surplus toward sellers. Bidders do not fully undo reporting bias and, when an ethical statement is present, learn more slowly about the seller's strategic behavior. Using a structural model of heterogeneous lying costs, we show that markets would be more price-efficient if bidders rationally internalized the distribution of lying costs. However, disclosure interventions aimed at enhancing ethical conduct need not reduce bias and can disadvantage bidders. Instead, they introduce 'ethical noise': shifts in beliefs without corresponding shifts in behavior that distort market outcomes.

Keywords: Misreporting, experiment, ethics, capital markets.

JEL Codes: D53, G14, G41.

"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 workings of markets depend not only on supply, demand, and regulation, but also on the ethical conduct and disclosure practices of market participants (Smith 1776; Stevens 2023). Because most statements cannot be verified, honest communication becomes a critical production input for society, governing everything from scientific integrity to financial disclosure (Marco and Larkin 2000; Yang 2009; Kristal et al. 2020a). Ethics, in this sense, is not a soft add-on to economics, it is a core input into the production of trust, which underpins all market exchange (Arrow 2006). Whether traders explicitly disclose an ethical stance, or remain silent about it, can therefore shape how markets process information.

Despite broad agreement that ethics matter for markets, our theoretical models capture only part of this connection. In the standard framework, if the incentives to misreport are common knowledge, traders discount reports accordingly, and a rational expectations equilibrium prevails: biased messages are anticipated and no one is misled (Dye 1988; Stein 1989). Actual communication rarely fits this benchmark. People routinely inflate résumés or embellish products, and counterparties only imperfectly undo these distortions, leaving systematic bias and added noise. Recent models introduce noise explicitly (Fischer and Verrecchia 2000; Dye and Sridhar 2004; Frankel and Kartik 2019), but we still lack a clear account of how noise, strategic reporting, and heterogeneous ethical preferences interact to shape market outcomes.

We address this gap using a controlled market experiment[1] in which a privately informed seller auctions an asset to competing bidders. Communication in this setting is potentially distorted by

[1]The experimental design has been pre-registered and is available for review at Open Science Framework: https://osf.io/ur943.

the costs associated with misreporting, creating an environment in which to explore the trade-offs between honesty and strategic misrepresentation. Specifically, we address three primary research questions: (1) Can bidders, given an understanding of the seller's economic incentives, effectively undo the seller's reporting bias? (2) To what extent does heterogeneity in aversion to lying explain the 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 anchored in the foundational principles of well-functioning markets, with roots tracing back to Adam Smith's 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 classical perspective finds resonance in contemporary analyses of trust dynamics such as the tragedy of the commons and the prisoner's dilemma (Kreps et al. 1982), and extends to empirical explorations of social capital, investor trust, and the regulation of financial professionals (Barrios et al. 2021; Guiso et al. 2008a; Zingales 2015; Kowaleski et al. 2020). Building on these foundations, we argue that the interplay between ethics, trust, and trade is more complex than commonly perceived. The heterogeneity in ethical preferences across individuals, combined with the inherent challenges of directly observing these preferences, gives rise to a phenomenon that we term "ethical noise." Ethical noise confounds the process of accurately inferring predictable biases, in stark contrast to the theoretical ideal of perfect communication under rational profit-maximizing behavior (Dye 1988; Stein 1989). Perhaps counter-intuitively, prices may be more efficient if it is common-knowledge that sellers exhibit no ethical aversion for lying.

Our empirical approach employs a controlled laboratory implementation of the classic single-period model of strategic communication, incorporating lying costs (Guttman et al. 2006; Kartik 2009; Kartik et al. 2007). This model has been extensively referenced in the literature on strategic communication and misreporting (Dye 1988; Stein 1989; Einhorn and Ziv 2008; Bertomeu et al. 2013; Caskey and Laux 2013). In our experimental setup, a seller, having privately assessed the value of an asset, makes a public report prior to its sale. The cost associated with reporting

increases proportionally to the degree of bias in the information provided, introducing a trade-off between strategic misrepresentation and the associated costs. To explore the role of ethical signaling, we include a separate treatment condition where sellers are mandated to sign a public ethical statement before proceeding with the sale. The monetary payoffs for the sellers are contingent upon both the extent of the bias in their report and the winning bid. Bidders, 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. Bidders' monetary payoffs are linked to the accuracy of their value prediction and the net profit from winning the auction. To ensure mutual understanding between sellers and bidders, participant roles are rotated in each round, with individuals being randomly rematched into groups of five from the experimental pool.

A potential concern is whether our use of a first-price auction mechanism affects the inferences we draw about the impact of ethical statements on market outcomes. This concern is unfounded. Our identification strategy relies on comparing outcomes between treated groups (where sellers sign ethical statements) and control groups (where they do not), holding the auction mechanism constant across both conditions. Since the first-price format is identical in both treatments, any effect of the auction mechanism on bidding behavior, price formation, or inference errors operates equally in treated and control groups. When we difference outcomes between the two groups, these common effects cancel out, leaving only the causal effect of the ethical intervention. The auction mechanism is thus orthogonal to our treatment effect as it affects the level of outcomes but not the difference between treated and control conditions. This differencing approach is robust to any systematic biases introduced by the first-price format, whether they stem from strategic bidding behavior, winner's curse effects, or other auction-specific dynamics.

Our experiment documents several notable results. First, bidders were generally unable to fully correct for the reporting bias of sellers. On average, the prediction error for winning bidders, measured as the difference between the true value of the asset and the value estimated by the bidders, was approximately \$3.40 (experimental dollars), while the asset's average fundamental value was approximately \$48.70 across all rounds of the experiment. Interestingly, when considering the col-

lective average estimate of all bidders, including those who did not win the auction, the deviation from the true value was even smaller (at just $1.27). 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 ability of the crowd to average out *strategic* errors, which, in our experiment, did not arise from the aggregation of private signals about the asset's true value (van Dolder and van den Assem 2018; Da and Huang 2020).

This finding contrasts sharply with predictions from rational expectations models that assume strict monetary payoff maximization. Such models assume that bidders, fully understanding the equilibrium strategy of the seller, can accurately infer the seller's private information and adjust their bids to reflect the true asset value, net of the equilibrium bias embedded in the report. Our results reveal substantial deviations from this theoretical benchmark. In the baseline treatment, variability in individual prediction errors led to significant heterogeneity in bidding behavior. Most notably, winning bidders overbid by an average of $11.87, allowing sellers to systematically profit from their biased reports.

Second, we observe the varied inference processes resulting from differences in seller preferences for honesty. Consistent with 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 aligned with payoff maximization. This bimodal pattern highlights the uncertainty bidders face in anticipating the ethical disposition of sellers (Marinovic 2013; Aghamolla and Smith 2023; Jiang et al. 2023). When we perform the same exercise to analyze bidders' beliefs based on their estimated biases, we find a more diffuse distribution. This diversity in beliefs translates into a disproportionate number of auction winners being bidders who assumed the seller was honest. In a counterfactual analysis under a rational expectations equilibrium framework, where bidders accurately infer the distribution of seller preferences and sellers act accordingly, ethical noise would have resulted in average inference errors of $6.3, compared to the $8.9 observed in the experiment. Even when accounting for changes in reporting strategies

due to common knowledge of rationality, our findings thus far suggests that rational play increases communication but also increases the equilibrium bias of the seller.

Next, we examine the impact of making ethics more salient in the auction by introducing experimental treatments that assess the influence of ethical codes on market behavior (Davidson and Stevens 2013). In these treatments, the sellers were required to sign an ethical statement before participating in the auction ("mandatory"). Contrary to expectations and the empirical findings in Heese et al. (2023), the introduction of ethical statements did not significantly alter the reporting behavior of sellers, irrespective of the ethical declaration.[2] A key feature of our experimental design, which is not the focus of Heese et al. (2023), is the separation of personal ethical concerns from reputational concerns (others' perceptions). In our study, neither bidders nor sellers are aware of each other's identities. Despite the limited impact of the intervention on sellers, we found a significant effect on bidders' behavior. The bidders perceived the ethical statement to imply that the sellers were more honest, which led to greater forecast errors and overbidding. Furthermore, bidders' inability to fully correct for the reporting bias appears to stem from a disruption in their learning process when they observe that the seller signed the ethical treatment. Overall, these results indicate that ethical interventions may have a more substantial impact on bidders, potentially leading to increased mispricing and ethical statement inadvertently shifting terms of trade toward sellers.

Our study offers several 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 incomplete or unenforceable. Motivating our study is the work of Sapienza et al. (2013), who show that global trust surveys often measure perceptions of others' honesty rather than individuals' own reporting behaviors. By examining the two sides of

[2]We recognize that this could suggest that our experimental setup might not have had sufficient power to induce notable shifts in honesty. However, our aim was not to design a mechanism to enhance ethics but rather to observe the implications of a reasonably framed ethical statement on bidders and sellers, and we used for this purpose ethical statements from prior literature and those used in the real world. 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 bidders, suggesting that the effect of ethics intervention may truly be greater on bidders and not simply reflecting a lack of power in the design.

a market, our experiment uniquely allows us to disentangle these two aspects - beliefs about trust and actual honesty - within a common laboratory setting.

Moreover, our study aligns closely with a substantial body of work in institutional economics that illustrates how trust-promoting institutions 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. (2021). In this literature, trust serves as a foundational element for well-functioning markets, particularly when contracts are incomplete or unenforceable. Within this broader literature, three particular studies share a thematic connection with ours, focusing on the credibility of reports when there is a potential misalignment. 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, which generally increases the likelihood of a sale. Finally, Sheremeta and Shields (2013) conduct a cheap talk experiment and find a tendency for subjects to expect more honesty from others than they experienced similar roles. While these studies focus on cheap talk, they collectively highlight the prevailing expectation for receivers to expect greater honesty than is warranted. 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.[3]

Furthermore, our study contributes to the understanding of trust in a two-sided market setting that involves transactions between bidders and sellers (O'Brien and Srivastava 1991; Forsythe et al. 1999; Bloomfield and O'Hara 2000; Crockett et al. 2009; Bloomfield et al. 2009). In existing research, trust is usually 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

[3]Our study gains additional relevance in light of recent developments in the field of ethics research. In particular, the retracting of an ethical signature study published in the Proceedings of the National Academy of Sciences (Kristal et al. 2020b) 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.

where one player has the ability to 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 aspects. First, in our market setup, participants divide the gains from trade, this removing the additional consideration of trust games where joint production creates additional motives for cooperation. While cooperation on a joint production is an important part of economic institutions, it is not always present in any market-level exchange; for example, a sale of securities to a shareholder or a market transaction involving a sale of real goods does not necessarily involve a joint enterprise. To focus only on the specifics of such arm's length transactions, we have thus removed mutual cooperation motives, except in terms of a liquidity need to sell the asset. This distinction is meant to reflect a transactional nature of markets, where exchanges are often impersonal and outcomes are driven by individual strategic behavior rather than shared cooperative goals. Additionally, previous research suggests that market-mediated exchanges might actually erode moral sentiment, replacing ethical considerations with profit-oriented decision-making (Falk and Szech 2013). Second, trust in our experiment pertains to biases in reporting about a fundamental asset value, while in trust games with communication, it often relates to the credibility of promises regarding future cooperation. In summary, these distinctions position our study to capture dynamics specific to pre-trade communication.

Our research also builds on studies that examine the failures of rational expectations in scenarios involving verifiable disclosure, where a privately informed sender decides whether to 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. These failures are often, theoretically, attributed to communication frictions, such as non-verifiability or the costs associated with disclosure (Beyer et al. 2010).

Our setting incorporates a strategic communication game that requires high-level strategic thinking, presenting participants with significant challenges in processing information based on the actions of others (Corgnet et al. 2018). This setting diverges from the verifiable disclosure

literature in several conceptual dimensions. For example, lying by omission, as in withholding information in a disclosure setting, may differ fundamentally from misreporting. In other words, the sender is not making a factually untruthful assertion when choosing not to disclose unfavorable information and there is typically no moral expectation of full disclosure. This reduces the scope for ethical triggers relative to direct misreporting. Conversely, in scenarios involving costly misreporting, the sender's actions are designed to obscure the true signal strategically. Even in equilibrium, the receiver is not entirely deceived, as the true signal can still be inferred (Dye 1988; Stein 1989). In contrast, under truthful voluntary disclosure, withholding reduces the overall informativeness of the communication, altering the dynamics of trust and strategic interactions.

Our analysis is organized along four sections. Section 2 presents the theoretical framework and empirical design, including the information known by bidders 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 bidders 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.) $F(.)$, 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.[4] 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 sellerâ€™s problem yields an interior solution, the optimal report r that meets the first-order condition in equation (1) is

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

[4] 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 bidders against overbidding.

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 bidders 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 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.[5]

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 bidder'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 bidder's estimate of the seller's true asset value.

[5] 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 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.

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.

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.[6]

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 bidders, 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.[7] 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.

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

[7] 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.

2020). 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 from a uniform distribution drawing values in $\{0, 1, \ldots, 100\}$. 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 $[0, 100]$ (window A.1), which we allow because, in the equilibrium in (4), a 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" treatment, sellers are required to certify the report and, 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 mainly explore two distinct treatments[8] in this regard: (1) no ethics statement, (2) mandatory ethics statement. In the ethics treatment (2), a refusal to sign results in the termination of trading for that session with the seller receiving no payoff and the asset remaining unsold. Bidders are informed about the experimental treatment in play.[9] 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), where the asset is sold in a first-price auction with four competing bidders. Bidders 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 payoff 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, bidders learn whether

[8]We also explore a voluntary treatment such that the seller may decline to sign, but still be allowed to play. The results regarding the effects of ethics statement closely mirror those from the mandatory treatment, so they are not included in the figures and tables for simplicity.

[9]It is a between-subjects design so that subjects are assigned to either the no ethics statement treatment or the mandatory ethics statement treatment.

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 trust in others, following the framework established by Davidson and Stevens (2013). These responses are collected in window A.5. The sessions were held online, with participants recruited from the subject pool of the Experimental Social Sciences laboratory at a prominent university.

2.3 Discussion and Objectives

Our experiment is designed to investigate a narrow subset of questions concerning trust and trade. While many additional factors may influence exchange behavior, we intentionally limit the scope of our inquiry to minimize confounding trade-offs and focus precisely on our research question. As a result, we exclude several potential research avenues, particularly those involving mechanisms already explored in prior literature. The experiment is constructed within a context chosen to align closely with the conceptual focus of the study, rather than to achieve descriptive realism. Below, we outline key design choices and justify them in relation to the scope of the study.

First-price auction mechanism. Our baseline experiment employs a first-price auction, where the transaction price corresponds to the highest submitted bid. This format is commonly used in experiments (e.g., Cox et al. (1988), King (1996), Goeree and Offerman (2002); see also the survey by Kagel and Levin (2011)). While alternative formats such as second-price or ascending auctions exist, we opt for the first-price format for several reasons. First, many participants lack formal training in game theory and may not recognize the dominant strategy in a second-price auction. Second, ascending and descending formats are less universal and introduce additional dynamics not central to our research.

Under standard rationality assumptions and absent private information, auction formats yield equivalent equilibrium outcomes, namely, the full dissipation of bidder surplus. However, when cognitive biases affect bidding behavior, the choice of format becomes consequential. The first-price format emphasizes heterogeneity in bidder assessments, as the winner's payment directly

influences terms of trade. Our presumption is that such heterogeneity is important in practice for outcomes in market exchanges. In contrast, second-price or ascending auctions tend to dampen, but do not eliminate, such heterogeneity, since the transaction price is set by the second most optimistic bidder. In counterfactual analyses, we simulate outcomes under a second-price auction using bidders' valuation reports (assuming dominant strategies), and find that this format shifts the terms of trade in favor of bidders.

Ethics statement design. Our study does not seek to develop or test targeted interventions aimed at promoting ethical behavior. Unlike research that examines the impact of explicit penalties, incentives, or psychological framing devices (e.g., Kristal et al. (2020b); Mazar et al. (2008)), we avoid engineered mechanisms that might artificially enhance honesty. Instead, we adopt a broad and realistic code of ethics modeled after Davidson and Stevens (2013), reflecting the kind of ethical declarations commonly encountered in market environments.

This design allows us to assess how market participants naturally respond to ethical prompts without introducing confounding experimental manipulations. We deliberately avoid framing effects or moral priming to maintain realism and focus on general behavioral tendencies rather than intervention efficacy. The ethical statement itself is borrowed from prior literature to ensure comparability and minimize researcher discretion. While stronger interventions could have been considered, we did not anticipate that the ethical statement would have no effect on average seller bias. Moreover, modifying the statement post hoc in response to the null result would compromise the integrity of the design. Importantly, our analysis focuses on differences in reported bias induced by the ethical intervention, rather than the overall existence of a treatment effect on sellers. Given ongoing debates in the literature about the effectiveness of ethical nudges, our findings should nevertheless be interpreted within the context of this specific experimental design and not generalized to all forms of ethical intervention.

Explicit versus Implicit Costs. A central aim of our study is to distinguish between explicit and implicit costs of misreporting, consistent with a substantial body of prior work. We implement explicit costs through quadratic penalties for sellers who misreport, capturing standard enforcement

mechanisms present in many institutional settings. Except in the knife-edge case of zero cost, such a structure yields a fully revealing equilibrium (e.g., Dye 1988; Stein 1989; Guttman et al. 2006). By contrast, with only implicit costs, understood as psychological discomfort or disutility from lying, there is no clear theoretical reason to expect that higher implicit costs would lead to more misreporting or reduced price efficiency. We thus hypothesize, following Gneezy et al. (2011), that participants consider both monetary and non-monetary consequences in their decisions.

To isolate these effects, we ensure that the ethical statement carries no monetary penalties. We also exclude other well-documented confounds, such as repeated interactions or reputational concerns. While repeated play has been shown to enhance cooperation and reporting accuracy (Schwartz et al. 2000), it is not possible in all real-world transactions and requires conditions (e.g., patience, coordination, or observability) that may not always be feasible. Our interest lies in mechanisms, like ethical prompts, that may foster improved communication and trust even in the absence of repeated interaction.

3 Reduced-Form Analysis

In this section, we present the reduced-form evidence that highlights the key behavioral patterns observed in our experimental setup. This analysis serves as a foundation for understanding how ethical interventions and strategic communication influence market outcomes, specifically focusing on the interactions between the sellers and bidders. By documenting these patterns, we provide empirical insights that inform the subsequent structural analyses and broader theoretical implications of our findings.

3.1 Descriptives

We ran three sessions for each of the treatments, (i) no ethical statement, and (ii) mandatory ethical statement. Table 3 provides descriptive evidence on participants' actions.[10] The mean asset value across both regimes is consistent (approximately $49), with minimal variation between

[10]In untabulated analysis, we also ran three sessions in which ethical statement is voluntary. The results closely mirror those from the mandatory treatment.

the "No Treatment" and "Ethics Treatment" groups. Also. the mean reports issued by sellers in the ethics treatment ($55.52) and that in no treatment ($52.38) are comparable, suggesting that ethical declarations might not change seller's reporting behavior. However, predicted values and bids show increases in the ethics treatment regime, with bids rising from a mean of $29.30 in the no-treatment group to $36.57 under the ethics treatment, indicating a potential influence of ethical reporting on bidder behavior. In addition, standard deviations for bids are higher in the ethics treatment, reflecting greater variability in bidder responses under this regime. These results hint at a potential shift in market dynamics driven by ethical declarations, with sellers potentially leveraging higher reports to influence bidding behaviors.

[Insert Table 3]

In Figure 1 we consider all treatments to plot two-by-two mean differences of the variables listed in Table 3, with representations of the corresponding 95% confidence interval. This figure suggests that sellers over-reported on average, bidders' predictions and bids were below the sellers' reports, bidders' 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).

[Insert Figure 1]

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 bidders' bids as a function of reported values in Panel (b), controlling for age, gender, and risk preferences. In Panel (a), sellers over-reported by $4.79 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 $21.27 on average. Contrary to actual seller behavior,

bidders exert an increasing discount for higher reports, and bid above the report for sufficiently low reports (i.e. for reports below $25).

[Insert Figure 2]

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, we found no evidence that signing an ethical statement affected sellers' reporting bias, relative to no ethical statement at any asset value.[11] The average seller bias is $4.11 under no ethical statement, and $5.65 under mandatory signing.

[Insert Figure 3]

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

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

where $Report$ is the dollar amount the seller announces regarding the asset value, $Ethics$ is an indicator variable for the ethics treatment. This variable is set to one if a seller is presented with an ethical statement to sign. $Controls$ is a vector of controls for the individual seller (age, gender, risk preference) as well as the asset value given to the seller. In some specifications, we include fixed effects for each round of play $Round\ FE$.

The results in Table 4 are consistent with Figure 3. Across all specifications in Columns 1-3, we fail to find a statistically significant difference in the reporting behavior of sellers exposed to the ethical treatment. Instead, the difference in reports between sellers exposed to the ethical treatment and those who were not exposed are economically small and not with the expected sign. In Column

[11] In untabulated analysis we conduct mean comparisons of sellers' reporting bias between mandatory and no ethical treatment, using a t-test. Rather than observing a reduction in bias when signing an ethics statement is mandated, we find a statistically significantly higher average reporting bias (at 5% level) in the mandatory ethical treatment setting compared to the setting without ethical signature.

3, the coefficient on $Ethics\ Statement$ indicates that sellers exposed to ethical statement report nearly \$1.85 higher on average.

Overall, the reduced form results suggest that the ethical treatment did little to change the reporting behavior of sellers.

[Insert Table 4]

3.3 Bidders and Ethics

Next, we focus on bidders' behavior, similarly partitioning the bids based on whether the bidders were exposed to the ethics treatment. Figure 4 plots the average bid for the bidders over the distribution of seller provided reports. The bin scatter controls for the age, gender, and risk preferences of the participants.

[Insert Figure 4]

Figure 4 suggests that bidders respond to the ethical treatment by choosing higher bids relative to the non-treated bidders. The average bidder placed a bid lower than the reported value by nearly \$15.76 under no ethical statement, \$9.61 under mandatory signing.[12] Nonetheless, the difference in the average bid between settings probably relates 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.

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

where Bid is the dollar amount the bidder proposes 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 5 presents the results from the regression in equation (7). Consis-

[12] In untabulated analysis we conduct two-by-two mean comparisons of bids between mandatory ethics treatment and no ethics treatment, using a t-test. We find that the mean differences is statistically significantly different from zero at the 1% level. However, this initial result can be driven by difference in amounts reported by the sellers. Therefore, we rely on the multivariate regressions results.

tent with Figure 4, Table 5 suggests a significant positive effect of observing an ethical statement on bids. This is true across all specifications.

In Column 3, where all the control variables are included, the coefficient on $Ethics\ Statement$ is statistically significant at 5% level, and translates to bids that are on average nearly $8.38 higher in the presence of ethics treatment. This is economically significant, given that the mean bid amount in the absence of treatment is $29.30. The results suggest that bidders may perceive an environment where signing an ethical statement is the norm as being less conducive to non-ethical behavior by sellers, compared to an environment where no signature of ethical statement is imposed.

[Insert Table 5]

Our finding that ethical statements increase the size of the bids is robust to potential issues from our first-price auction design. Specifically, by comparing ethical disclsoure treated bids to non-treated bids ensures that the first-price auction mechanism does not confound our treatment effect estimates, since the auction format is held constant across treated and control conditions. In Appendix 5, we provide an additional robustness check, excluding the winning bids from the analysis. Across all specifications, the positive effect of ethical statements on bidder behavior remains statistically and economically significant, reinforcing our conclusion that ethical declarations influence market outcomes primarily through their impact on bidder expectations rather than seller behavior. We also observe similar results when we examine the relation between ethics and trust on bidding behavior after excluding the winning bids.

3.3.1 Bidders Learning and Ethics

Although we ensured that subjects play the role of both bidders and sellers in different rounds, a concern that may arise is that participants may be more prone to behave naively due to misunderstanding of the economic environment, with the results likely capturing the heuristics they employ. Therefore, we assess whether, in general, bidders are able to learn and adjust their prior accordingly. More specifically, we assess how a bidder's prediction error evolves over the rounds

within each session, using the regression in equation (8).

$$Prediction\ Error \quad = \quad \beta_0 \times Round + \sum \gamma_k \times Controls_k + Subject \times Session\ FE + \epsilon, \quad (8)$$

where *Prediction Error* is measured as the absolute value of the difference between the true asset value and a bidder's estimate of the asset value. *Round* indicates the number of rounds played by a bidder during a session. Our assessment is based on the assumption that predictions merely guided by heuristics should not improve as more rounds get played (i.e. β_0 statistically insignificant) with and without the ethical treatment. Columns 1 and 2 of table 6, show the results from running the regression in equation (8) without and with ethical treatments, respectively. In column 1, the coefficient β_0 is negative and statistically significant at 1% level, indicating that for every additional round played, bidder's prediction error decreases by \$0.23, absent any influence of the treatment. This result suggests that the experiment setting itself does not prevent bidders from learning. However, in column 2, the coefficient β_0 is statistically insignificant, implying that the ethical treatment may have an effect on bidders' tendency to learn. To better capture the significance of the differential in learning between the regime with ethical treatment and that without treatment we run the regression specified in equation 9.

$$\begin{aligned} Prediction\ Error \quad = \quad & \beta_1 \times Round \times Ethics + \beta_2 \times Ethics + \beta_3 \times Round \\ & + \sum \gamma_k \times Control_k + Subject \times Session\ FE + \epsilon \end{aligned} \quad (9)$$

The differential in learning between the two regimes is captured in the coefficient β_1. The results from the regression in equation (9) are presented in column 3 of table 6. The coefficient β_1 is positive and statistically significant at 5% level, suggesting that the decrease in bidder's prediction error from one round to another is \$0.17 more when there is no ethical statement compared to when an ethical statement is signed by the seller.

[Insert Table 6]

The absence of learning in the ethical treatment setting may stem from increased reliance on

the seller's report when bidders see that the seller has signed an ethical statement. We explore this possibility further in the next section.

3.4 Trust and Ethics

We next 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 Column 3 of Table 4 for the sellers and of Table 5 for the bidders on the partition samples based on trust. Table 7 provides the results, with Columns 1-2 for sellers and 3-4 for bidders.

[Insert Table 7]

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 bidders, low-trust subjects choose higher bids in the ethical statement treatment. In contrast to high-trust bidders who exhibit low levels of skepticism regarding sellers' reports even absent ethical statement, the low-trust bidders become more confident in sellers' reports only when an ethical statement is signed. Specifically, the coefficient on *Ethics Treatment* in the low-trust samples is 12.815 which is economically meaningful, and statistically significant at 1% level, while the coefficient in the high-trust sub-sample is statistically insignificant and economically small.

4 Structural Estimation

Building on the insights from the reduced-form analysis, in this section, we employ structural estimation to rigorously quantify the underlying mechanisms driving market behavior in our experimental setup. By estimating key parameters, such as the distribution of ethical preferences and

the pricing function, we aim to disentangle the roles of strategic misreporting, bidder inference, and the impact of our ethical intervention.

4.1 The Seller's Problem

We begin by presenting structural evidence that sellers do not uniformly behave as profit maximizers (i.e., $e \neq 1$ for most types), and that there is significant heterogeneity and uncertainty in ethical preferences. To analyze this, we use equation (5) to estimate the pricing function $\hat{P}_i(.)$ non-parametrically, conditioning on whether there is an ethical statement ($i = 1$) or not ($i = 0$).[13] We focus on the winning bid as this represents the key payoff-relevant outcome for sellers.

As noted above, in many first-price auction experiments (and ours is no exception), participants often exhibit overbidding behavior (Goeree et al. 2002; Crawford and Iriberri 2007). This tendency may stem from strategic errors or from bidders deriving additional utility from winning beyond the monetary payoff. In the analysis that follows, we first take the sellers' perspective, without presuming that bidders adhere to equilibrium bidding strategies. For estimates relating to bidders, we use the bidder's induced valuations of the seller's report to infer bidders' beliefs, separately from their observed bidding behavior. Nevertheless, it is important to note that overbidding disproportionately penalizes the winning bidders, and this effect intensifies with greater heterogeneity in trust levels among participants. To maintain the robustness of our analysis, we drop three observations where sellers reported values below zero. These anomalous reports likely reflect exploration or learning during the experiment and deviate significantly from the expected reporting strategies.[14]

[Insert Figure 5]

Figure 5 plots the estimated selling price, comparing treatments with versus without the ethical

[13] We use the optimal cross-validated bandwidth with the Epachnikov kernel. Confidence intervals were obtained by bootstrapping with 300 repetitions. We opted for a kernel-based method over partitions (as in the reduced-form model) because the slope of the pricing function needs to be estimated.

[14] 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.

statement. As in the reduced-form analysis, we also verify that bidders become more skeptical over higher reports, and signing an ethical statement reduces bidder skepticism.

[Insert Figure 6]

In Figure 6, we plot the histogram of inferred ethical types e using (5) for each seller report, using the difference between report and true signal (which we observe in the experiment), and the estimated $\hat{P}_i'(.)$ in Figure 5. These distributions are similar.[15] Both distributions are bimodal with truthful reporting ($e = 0$) representing the most common type, and there is a second mode located roughly at the payoff-maximizing strategy ($e = 1$).[16] That many observations feature seemingly sub-optimal play $e \notin [0, 1]$ may also reflect noise. To average out the 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 (40%) of all subjects preferring truthful reporting (defined as $e \in (-0.5, 0.5)$) and 32% with behavior roughly consistent with maximizing payoffs ($e \in [.5, 1.5]$). The median (average) type is $.65$ ($.67$), which falls somewhere in-between truthful reporting and maximizing economic incentives. Approximately 82% have an estimated type $e \in [-0.5, 2]$), implying that some of the heterogeneity at the observation level with high levels of misreporting may be partly due to noise.

[Insert Figure 7]

Table 8 documents differences between seller types, under two specifications, (1) ordinary least squares and (2) random effects. 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

[15] In untabulated analyses, we conduct a Kolmogorov-Smirnoff and Wilcoxon tests and find that both tests fail to find a significant difference between the distributions at conventional significance levels.

[16] There is some heterogeneity around these modes with, for example, some subjects choosing biases that are too large relative to the potential price benefit ($e > 1$). One explanation for heterogeneity among observations may be that some of the bidders 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).

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. The intercept is around two thirds across all specifications, similar to the average type in the histogram.

Consistent with the reduced-form analysis, we do not find a significant effect of signing an ethical statement, relative to the baseline of no ethical statement. The effect of the ethical statement is economically small and insignificant. We do not find that repeated play affects ethical types as, after excluding the practice round, the effect of the number of rounds is insignificant and economically small. For example, the effect of 10 rounds on e is roughly 0.1, which would map to roughly $\$1$ of incremental bias in equation (3) and is much less than the between-subject variation.

[Insert Table 8]

4.2 The Bidders' Perspective

Thus far, we have focused on sellers' reporting choices, considering the *actual* distribution of ethical types and how they perceive bidders. In what follows, we examine in more detail bidders' *perception* of sellers' types. We begin with a simple introductory question: are bidders able to accurately infer the seller's information, as would be predicted by the rational model? To answer this question, we asked the bidders in the experiment to provide their best assessment of the asset's value 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.

[Insert Figure 8]

We consider the bidders' prediction accuracy relative to actual bias behavior in the experiment, defined as the difference between the prediction and the true asset value. Note that bidders 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 bidders' prediction errors over each treatment. On average, absolute prediction errors are $7.45 and $9.6 in the

mandatory ethics and no ethics treatments, respectively. Overall, we find evidence of a "strategic" wisdom of the crowd effect such that, *on average*, the combined estimate of all bidders is able to reasonably correct the bias. Interestingly, because players in the experiment have no private 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 suggest that these functions may have a more important role in thin markets (Gerakos et al. 2013). However, 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. In untabulated mean difference tests, we find that the absolute errors are significantly lower (at a 1% confidence level) with an ethical statement; however, the reduction is small, with a differential of $2.15.

One explanation for this failure of rational inference may be that bidders 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 $7.3 and $7.1 in the mandatory ethics and no ethics treatments, respectively. This difference is no longer statistically significant, further suggesting that ethical statements do not appear to reduce prediction errors in economically meaningful magnitudes.

[Insert Figure 9]

Next, we compare the actual distribution of seller types to the inferred distribution as perceived by bidders. Using equation (5), we calculate the bias $r - \hat{v}$ reported by the bidder, i.e., where r represents the seller's reported value, and $\hat{v}$ is the bidder's predicted value of the asset. In Figure 9, we observe that the distribution of inferred ethical preferences closely resembles the distribution of actual seller types. As in the seller data, we find a significant mass point corresponding to honest sellers, but the peak around profit-maximization ($e = 1$) is less pronounced. This disparity may arise because the process of inferring ethical preferences indirectly, based on observed reporting

and bidding behavior, is more complex for bidders than for sellers who act directly based on their own preferences.

Interestingly, this result is also somewhat counter-intuitive. One might expect the inferred distribution of ethics to be *less* dispersed than the actual distribution since bidders' beliefs, while noisy, should converge toward average seller behavior. Instead, our findings reveal the opposite: we observe that bidders exhibit significant disagreement in their perceptions of seller ethics, as reflected in the broader dispersion of inferred types in Figure 9. This discrepancy highlights a critical source of ethical noise: it is not solely driven by heterogeneity in actual seller ethics, but also by the divergence in bidder perceptions of these ethical difference. Bidders' limited common knowledge about the distribution of seller preferences exacerbates this disagreement. In other words, neither sellers nor bidders appear to have accurate higher-order beliefs about what other parties believe about ethics. Nevertheless, the average perceived ethical type aligns closely with the actual average type, as shown in Figure 6. This suggests that while individual bidders struggle to infer the precise ethical preferences of sellers, the aggregate perception is relatively accurate. The persistence of this pattern underscores the nuanced interaction between ethical heterogeneity and market inefficiencies, where both actual behavior and perceived ethics jointly shape the market outcomes.

4.3 Counterfactuals: Rational Expectations and Second-Price Auctions

Up to this point, our analysis has considered sellers and bidders without assuming that either group necessarily maximizes their (ethical) preferences under the hypothetical condition of having knowledge of the distribution of ethical types. In this section, we examine the potential benefits and unintended consequences of increasing awareness about ethical noise, aiming to facilitate the convergence of bidders' and sellers' beliefs toward greater accuracy. Specifically, we ask how a level of skepticism appropriately calibrated to actual preferences can help reduce ethical noise and improve communication. Note that an answer to this question is a-priori unclear, as greater rationality may lead to more accurate bidder beliefs, holding the seller's actions constant, but the

resulting adjustments in pricing dynamics 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 sellers and assuming that both bidders and sellers choose mutual best response strategies.[17] Unlike classical setups where true values and ethical preferences follow standard distributions (e.g., a normal), the model precludes a closed-form solution for the estimated empirical distribution of ethical types. To address this complexity, we develop and implement 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 $e \in E = \{e_1, \ldots, e_J\}$, yielding $101 \times J$ sample points (v_i, e_i). Starting with an initial price function $P_0(r) \equiv r$, we solve for the seller's best response report $R_0(v_i, e_i)$ using P_0 in equation(1) and finding the global maximum over reports $r \in \{0, 1, \ldots, 120\}$.[18] The price function is then updated iteratively using a learning rate ρ to incorporate the expected value of the asset conditional on the observed report:

$$P_1(r) \equiv \rho \mathbb{E}(\tilde{v}|R_k = r) + (1 - \rho)P_0(r), \tag{10}$$

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 (10), and updating P_{k+1} until the average difference $|P_k - \mathbb{E}(\tilde{v}|R_k = r)|$ is small.

It may seem intuitive to set $\rho = 1$ in equation (10), which corresponds to iterated best responses, where the pricing function is repeatedly updated to the seller's strategy, and the seller's report is re-optimized to match the new pricing function. However, numerically, this algorithm does not converge. For two given reports close to each other, r_1 and r_2, which are not yet aligned with the equilibrium pricing function, the price $P_k(r_2)$ may temporarily be much higher than $P_k(r_1)$. This discrepancy encourages unethical sellers to shift toward r_2. In the next iteration, the report r_2

[17] We use the median e by subject, as shown in Figure 7, and exclude $e < 0$ to reduce estimation noise and outlier effects.

[18] To facilitate comparison with the fully-separating equilibrium in (4), we allow for reports greater than 100. However, the analysis is very similar if we bound reports by 100 since only values very close to 100 will attain this upper bound.

will be heavily discounted, making $P_{k+1}(r_1)$ appear relatively attractive as compared to $P_{k+1}(r_2)$. This cyclical behavior causes the algorithm to oscillate without stabilizing. By reducing the learning rate (ρ), the algorithm updates the pricing function more gradually, avoiding such oscillations. We find that setting $\rho = 0.001$ ensures convergence, allowing the equilibrium conditions to be verified with a high level of precision.

In Figure 10, we plot the equilibrium reporting function. This function features a degree of calibrated skepticism by bidders, typically midway between truthful reporting and payoff-maximization. As reports become more aggressive, the pricing function converges toward a baseline scenario assuming no ethical preferences, where higher reports are viewed as more likely originating from less ethical sellers. This behavior demonstrates how bidders dynamically adjust their strategies, assigning increasing skepticism to higher reports and incorporating the expected bias from sellers with weaker ethical preferences.

[Insert Figure 10]

In Table 9, we compare experimental outcomes to those predicted under the rational expectations equilibrium, focusing on pricing errors, bias, and seller welfare. Note that one can infer from this table the bidders' surplus by subtracting the average pricing error, that is, a higher average pricing error tends to move the terms of trade toward the seller.

[Insert Table 9]

Notably, subjects in the experiment have higher absolute pricing errors, i.e., lower price efficiency, than what they would have achieved under rational expectations even though, in equilibrium, they 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 roughly equal to the true value in the Nash equilibrium (Column 1) but it remains very small even in the experiment at $\$1.29$. However, the absolute pricing error is reduced by about a third if bidders follow the Nash equilibrium, showing that "rational" bids make price more informative by correcting extreme over

and under pricing. This occurs even though the average bias almost doubles from $\$4.94$ to $\$8.57$ (Column 3), as sellers adapt to greater skepticism by biasing more. In summary, this illustrates the trade-off between noise reduction and predictable bias: increased rationality among bidders reduces random pricing errors but amplifies the systematic bias as sellers adapt to tougher pricing discipline.

Sellers who prioritize maximizing their cash payouts ($e = 1$) are worse-off when bidders adopt more rational strategies. On average, these sellers earn $\$8.24$ less under rational expectations than in the experimental setting, as they are less able to exploit over-bidding. Compared to a scenario where all sellers maximize their payouts ("Nash cash maximizers"), these strategic sellers experience even greater losses, earning an additional $\$3.15$ less. The inability to camouflage as truthful sellers under greater bidder skepticism limits their ability to extract surplus.

Finally, we examine the effects of increased misreporting costs, doubling the penalty for biased reporting. We find that higher misreporting costs significantly reduce both predictable bias and pricing errors, with both measures approximately halved. Sellers are on average better-off, although not sellers who maximize cash flows ($e = 1$). This finding aligns with theoretical results by Fischer and Verrecchia (2000), who demonstrate that when preferences are uncertain, greater reporting flexibility may, under certain conditions, benefit sellers even though bidders price-protect, as it allows sellers to better adjust their reporting strategy as a function of their ethical preferences e. However, our results suggest a nuanced dynamic: while sellers with greater misreporting costs face stronger constraints, the reduction in bidder skepticism offsets this effect, yielding modest welfare improvements for sellers overall.

In columns 4 and 5, sellers are (marginally) better-off when misreporting costs are higher. This improvement is particularly notable even in the absence of uncertainty about ethical preferences, suggesting that the benefits of reduced pricing errors and bias are not tied solely to scenarios with uncertain preferences. Higher misreporting costs reduce the noise in bidder inferences, enabling bidders to better differentiate between sellers based on their reports. This, in turn, improves market efficiency by aligning reported values more closely with true asset values and reducing the infor-

mational asymmetry that often distorts transactions. Thus, the overall welfare effects of higher misreporting costs are nuanced, reflecting both the benefits of increased accuracy and the limitations imposed on strategic adaptability.

Next, we consider a second counter-factual exercise, altering the auction mechanism to a second-price (or descending) auction. While, in principle, the experimental design would allow us to run another experiment with a different selling mechanism, our inference of valuations by buyers allows us to form some conjectures of outcomes in a second-price price auctions. To solve this alternative auction and keep the analysis as straightforward as possible, we further assume that bidders can learn that the dominant strategy is to bid one's assessment; of course, in practice, it is possible that bidders would respond very differently but assessing this would require a formal experiment rather than a counter-factual inferred from first-price auctions.

Holding the seller's bidding constant, we recompute the terms of trade assuming that bidders (now) choose their valuation and the second highest price is selected. Table 7 reveals that because the resulting bids are smaller, the second-price format benefits the seller by reducing overbidding by the more aggressive (winning) bidder. Note that the second-price auction coincides with the Nash solution of the first-price auction if both sellers and buyers fully re-optimize their strategy, since in this case all (competing) bidders, who have the same information, bid the same expected value.

5 Conclusion

This study investigates the relation between ethical declarations and market transactions. Our results challenge the conventional view that enhanced ethical reporting inherently benefits markets. Our findings reveal that ethical considerations significantly influence participant behavior, regardless of whether an ethical statement is signed. This highlights the importance of psychological costs of lying as a driver of behavior. However, contrary to popular beliefs, standard ethical interventions had minimal impact on altering seller behavior beyond a baseline ethical conduct. Ethical statements only slightly modified reports.

Interestingly, the introduction of ethical statements had unintended consequences. Certain trusting bidders perceived these declarations as credible, leading them to overpay in response to sellers' reports. This phenomenon, which we call *"ethical noise,"* complicates price efficiency by introducing variability in bidder expectations. Unlike predictable biases, ethical noise reflects the difficulty bidders face in forming accurate beliefs about the ethical preferences of sellers, ultimately disadvantaging bidders. In markets where sellers are purely profit-maximizing, bidders may actually experience greater price efficiency due to the absence of this noise.

Although 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 bidders (via 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. Although effective in extracting preferences, this experimental design is confined to the laboratory setting and might not capture all considerations in real-world ethical decision-making, where factors such as reputational concerns, long-term relationships, and broader social dynamics play a critical role.

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 (2023) 'Strategic complexity in disclosure.' *Journal of Accounting and Economics* 76(2-3), 101635

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 (2020) 'Incentives in experiments with objective lotteries.' *Experimental Economics* 23, 1–29

Barrios, John M, Efraim Benmelech, Yael V Hochberg, Paola Sapienza, and Luigi Zingales (2021) '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 (2000) 'Can transparent markets survive?' *Journal of financial Economics* 55(3), 425–459

Bloomfield, Robert, Maureen O'hara, 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.' *Stan. L. Rev.* 56, 1

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

Cox, James C, Vernon L Smith, and James M Walker (1988) 'Theory and individual behavior of first-price auctions.' *Journal of Risk and uncertainty* 1, 61–99

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

Dye, Ronald A., and Sri Sridhar (2004) 'Reliability-relevance trade-offs and the efficiency of aggregation.' *Journal of Accounting Research* 42(1), 51–88

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, R, R Lundholm, and T Rietz (1999) 'Cheap talk, fraud, and adverse selection in financial markets: some experimental evidence.' *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

Gneezy, Uri, Stephan Meier, and Pedro Rey-Biel (2011) 'When and why incentives (don't) work to modify behavior.' *Journal of economic perspectives* 25(4), 191–210

Goeree, Jacob K, and Theo Offerman (2002) 'Efficiency in auctions with private and common values: An experimental study.' *American Economic Review* 92(3), 625–643

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

Heese, Jonas, Gerardo Pérez-Cavazos, and Caspar David Peter (2023) 'When executives pledge integrity: The effect of the accountant?s oath on firms? financial reporting.' *The Accounting Review* 98(7), 261–288

Hurkens, Sjaak, and Navin Kartik (2009) 'Would i lie to you? on social preferences and lying aversion.' *Experimental Economics* 12, 180–192. 10.1007/s10683-008-9208-2

Jiang, Xu, Baohua Xin, and Yan Xiong (2023) 'The value of mandatory certification: A real effects perspective.' *Journal of Accounting Research* 61(1), 377–413

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), 141–73

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

Kagel, John H, and Dan Levin (2011) 'Auctions: A survey of experimental research, 1995–2010.' *Handbook of experimental economics* 2, 563–637

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

Kowaleski, Zachary T, Andrew G Sutherland, and Felix W Vetter (2020) 'Can ethics be taught? evidence from securities exams and investment adviser misconduct.' *Journal of Financial Economics* 138(1), 159–175

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 (2020a) 'Signing at the beginning versus at the end does not decrease dishonesty.' *Proceedings of the National Academy of Sciences* 117(13), 7103–7107

_ (2020b) '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

Schwartz, Steven T, Richard A Young, and Kristina Zvinakis (2000) 'Reputation without repeated interaction: A role for public disclosures.' *Review of Accounting Studies* 5, 351–375

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

Zingales, Luigi (2015) 'Presidential address: Does finance benefit society?' *The Journal of Finance* 70(4), 1327–1363

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:

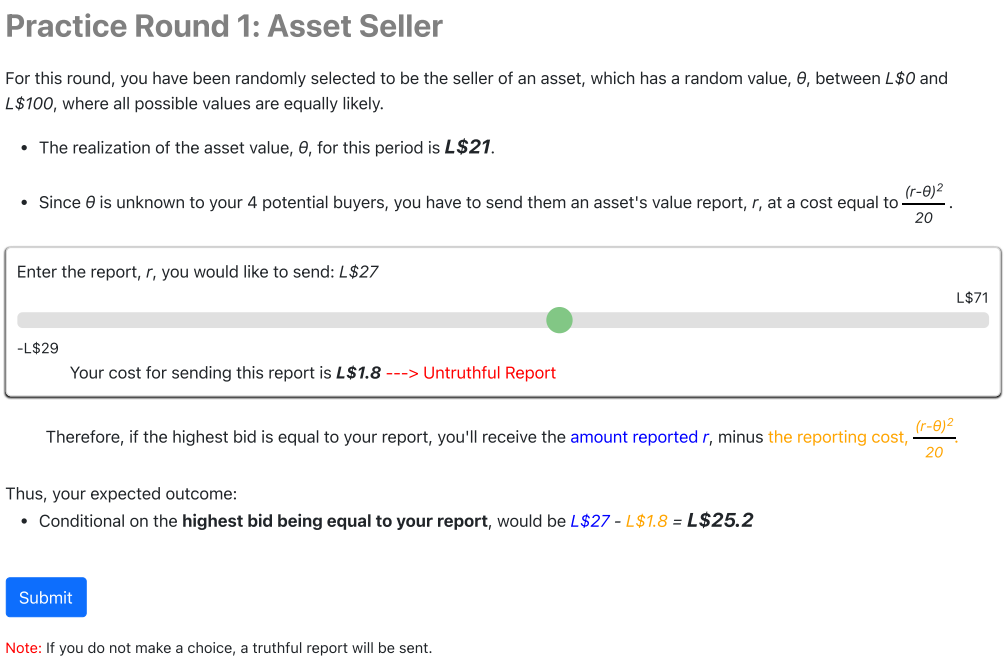

Note:
1. The subject can move the slider around to see what his reporting cost will be, plus the expected outcome conditional on the highest bid being equal to their report.
2. If the subject does not make a choice, the report defaults to the asset value.

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

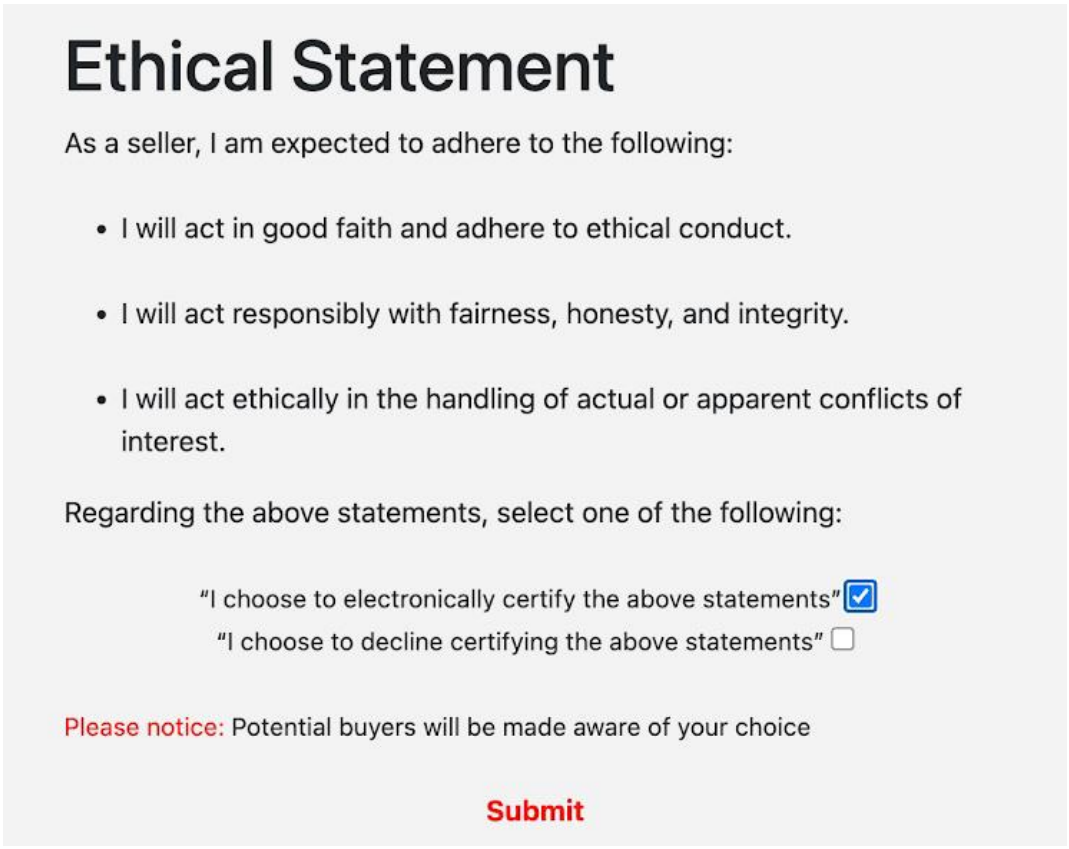

Window A.3 - bidder's Screen:

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 - bidder'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

Robustness Checks: First-Price Auction Structure

Table 1: Robustness: Regression Analysis of Bids Excluding First-Price Winners

	Bid		
	(1)	(2)	(3)
Ethics Treatment	9.751***	9.498***	8.273**
	(3.555)	(3.590)	(3.627)
Observations	1,277	1,277	1,277
Adjusted R-squared	0.029	0.045	0.250
Controls	No	No	Yes
Round FE	No	Yes	Yes

Notes - This table excludes winning bids of the first price auction, and summarizes the OLS regressions results of the amounts of bids by the other bidders on ethical treatment. The regressions are based on the sample of bidders only. The baseline is comprised of the bids in the sessions with no ethical treatment. *Bid* is the dollar amount the bidder proposes. *Ethics Treatment* is an indicator variable that takes the value of 1 when sellers are provided with an ethical statement they have to sign (bidders observe the signed ethical statement, if any); 0 otherwise. The control variables are comprised of individual seller characteristics (age, gender, GPA, and risk preference) as well as the true asset value in dollar amount. Round FE indicates whether the estimates are within rounds. The standard errors are clustered at the bidder level, and reported in parentheses below the coefficients. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Table 2: Robustness: Relation between Ethics and Trust, Excluding First-Price Winners

	Report (Sellers)		*Bid (Buyers)*	
	(1)	(2)	(3)	(4)
Ethics Treatment	1.637	1.731	9.998**	6.358
	(1.246)	(1.743)	(4.558)	(5.817)
Observations	221	217	666	611
Adjusted R-squared	0.914	0.902	0.301	0.188
Controls	Yes	Yes	Yes	Yes
Round FE	Yes	Yes	Yes	Yes
Subsample	Low Trust	High Trust	Low Trust	High Trust

Notes - This table excludes winning bids of the first price auction for the sample of buyers, and summarizes the OLS regressions cross-sectional results of the sellers' reports and bidders' bids on ethical treatment, when splitting each of the samples (bidders and sellers) 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. *Ethics Treatment* is an indicator variable that takes the value of 1 if the seller is presented with an ethics statement to sign (mandatory treatment); zero otherwise. Controls is a vector comprised of individual seller characteristics (age, gender, GPA, and risk preference) as well as the true asset value (for sellers sub-samples) or Value reported value (for bidders sub-samples). 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.

Figures

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

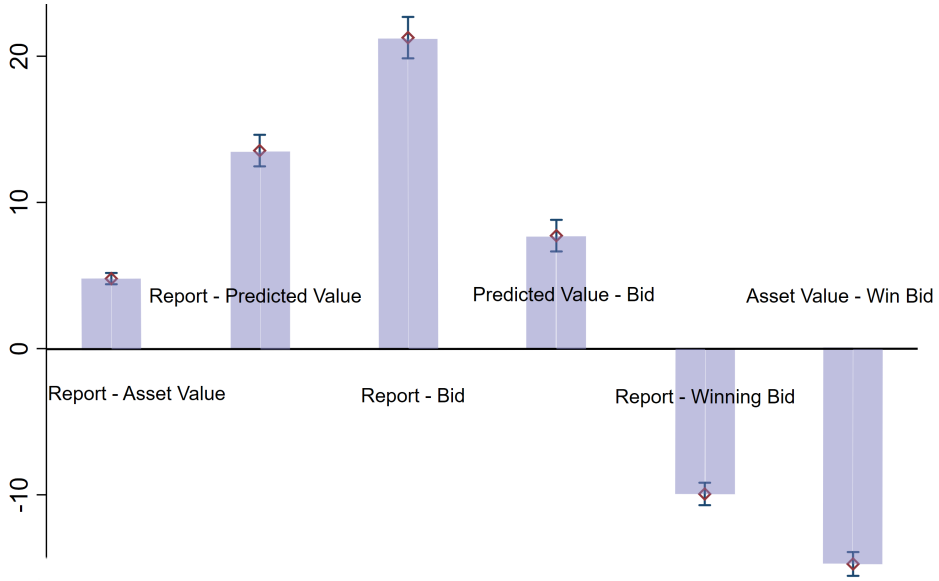

Notes - This figure shows the mean differences in diagram bars (with 95% confidence interval), across all treatment regimes, between : reported and asset value; reported and predicted value; reported value and bid; reported value and winning bid; predicted value and bid; asset value and winning bid.

Figure 2: Sellers Reports and bidders Bids

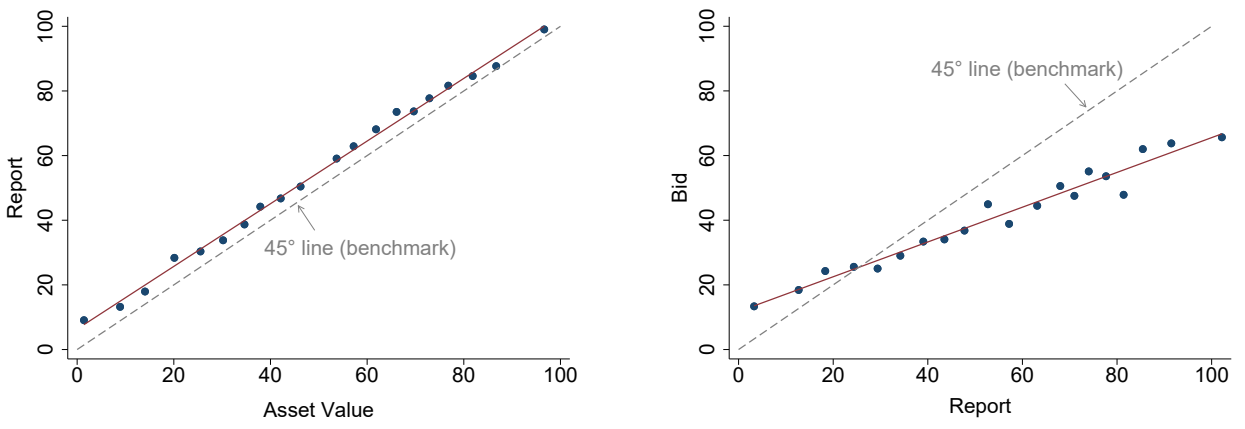

(a) Reported Value as a Function of Asset Value

(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 bidders' bids as function of sellers' reports. The 45-degree line represents truthful reporting in Panel (a), and bidders bidding the reported value in Panel (b).

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

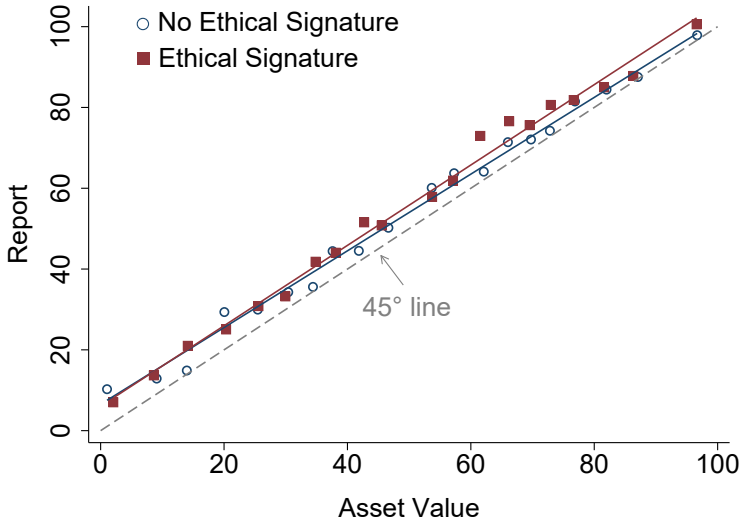

Notes - This figure presents binned scatter plots of sellers' reports as function of asset values. The annotation "Ethical Signature" represents the sessions where all sellers have to sign an ethical signature. The 45-degree line represents truthful reporting.

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

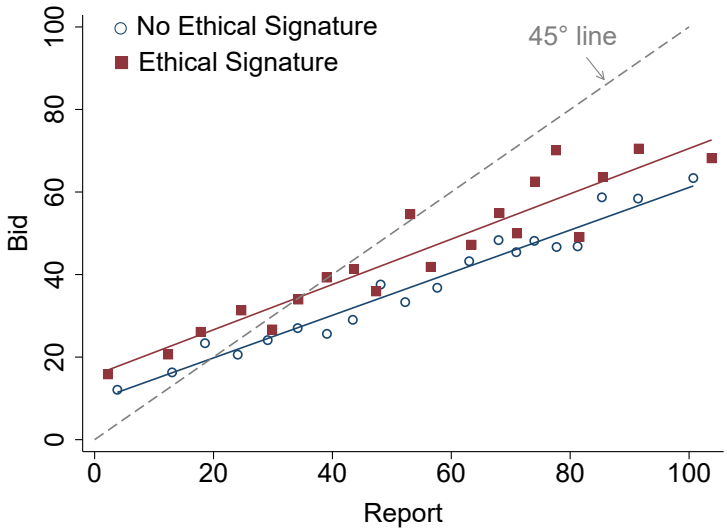

Notes - This figure presents binned scatter plots of bidders' bids as function of sellers' reports. The annotation "Ethical Signature" represents the sessions where bidders observe that sellers signed an ethical statement. The 45-degree line represents bidders bidding the reported value.

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

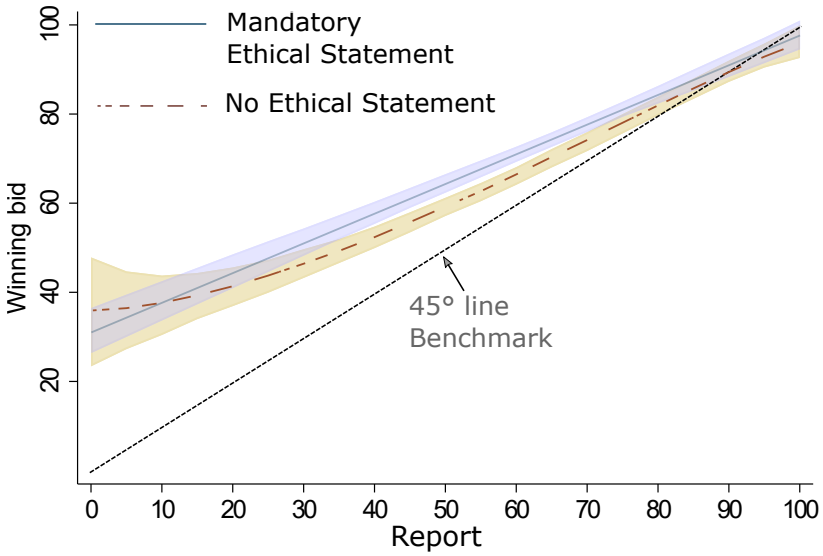

Notes - This figure shows the estimated selling price conditional on asset value reported in the mandatory vs no ethical statement settings. The 45-degree line represents bidders bidding the reported value.

Figure 6: Inferred Seller Ethical types

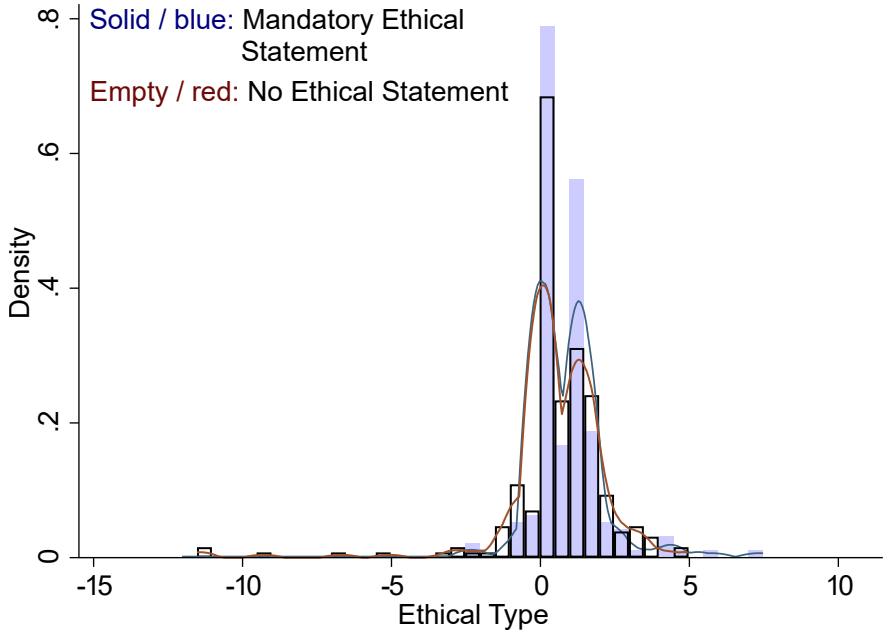

Notes - This figure presents the distributions of inferred ethical coefficients over Mandatory vs. No Ethical Treatment, as obtained from (5). The curve is a kernel fit of each distribution.

Figure 7: Ethical types averaged by subject

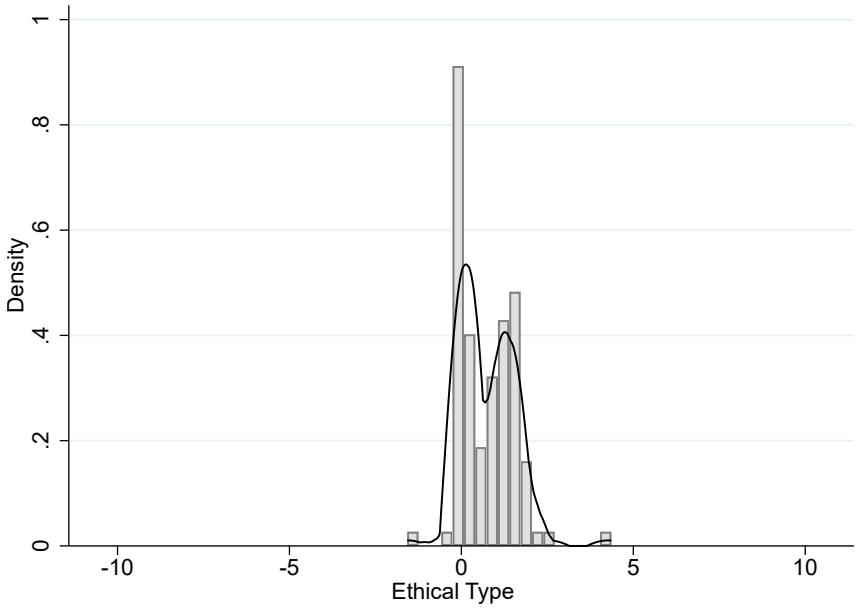

Notes - The distribution shown in this figure is obtained after computing for each subject the median inferred ethical type from (5).

Figure 8: Comparison of Prediction Errors

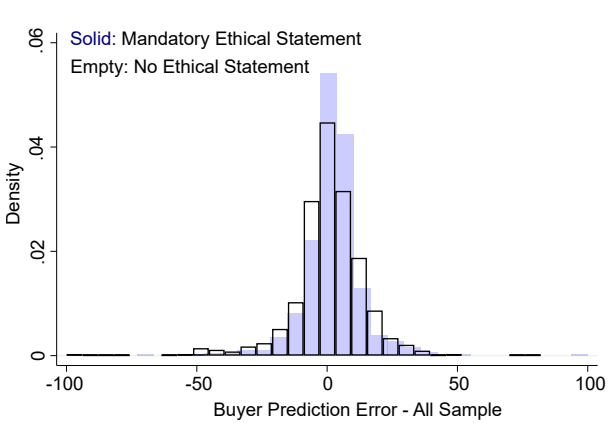

(a) Mandatory vs No Ethical Statement

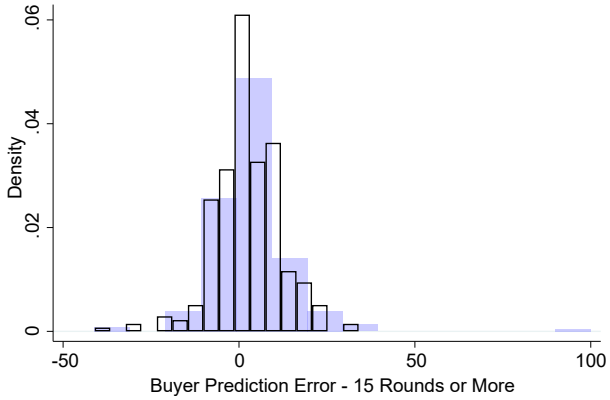

(b) Mandatory vs No Ethical Statement (Experienced bidders)

Notes - This figure shows the distribution of bidders' prediction errors in different treatments. Panel (a) includes all bidders under the mandatory vs no ethical statement treatment, while Panel (b) focuses on experienced bidders (those having played 15 rounds or more) to identify learning effects.

Figure 9: Ethical Types - bidders' Perception

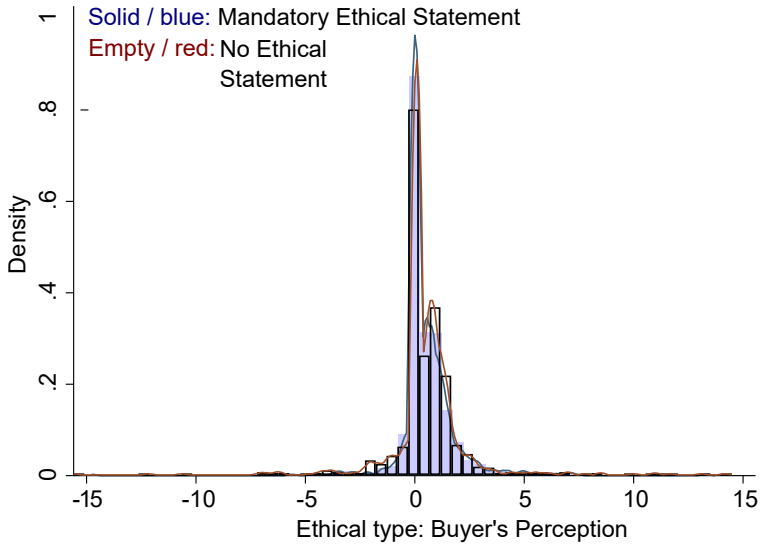

Notes - This figure shows the distribution of the estimated perception of bidders regarding sellers' types in the mandatory ethical statement and no ethical statement settings.

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

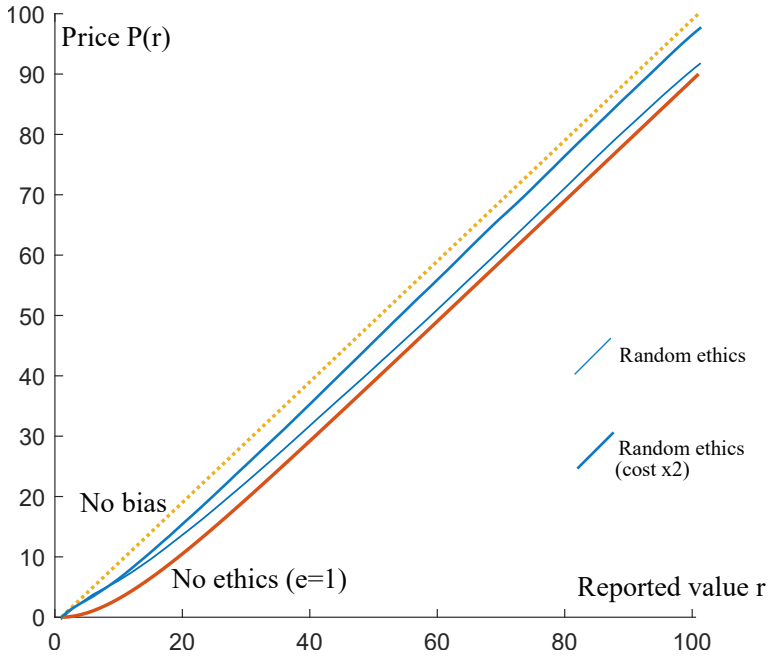

Notes - This figure compares the actual price function considering the reported asset value as input (curve referenced "Random ethics", generated from the results of the experiment) to the counter-factual price function (curve referenced "No ethics (e=1)" obtained by simulation from rational expectations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

Tables

Table 3: Summary statistics

Regime of Ethics Treatment	Variables	N	Mean	Std. Dev.	Min	Median	Max
No Treatment	*Asset Value*	1245	48.27	28.59	0	50	100
	Report	1245	52.38	29.30	0	55	110
	Predicted Value	1245	39.05	31.07	0	37	110
	Bid	1245	29.30	30.75	0	20	100
Ethics Treatment	*Asset Value*	969	49.86	28.55	0	50	100
	Report	969	55.52	29.30	0	56	115
	Predicted Value	969	41.71	31.96	0	40	115
	Bid	969	36.57	33.11	0	33	100
Both Regimes	*Asset Value*	2214	48.97	28.58	0	50	100
	Report	2214	53.76	29.33	0	55	115
	Predicted Value	2214	40.21	31.48	0	39	115
	Bid	2214	32.48	32.00	0	25	100

Notes - This table provides the summary statistics from each of the ethical treatment regimes. *Asset Value* is the value in dollar amount observed by the seller regarding the asset. *Report* is the dollar amount the seller announces regarding the asset value. *Predicted Value* is bidder's estimate of the true asset value. *Bid* is the dollar amount the bidder proposes.

Table 4: Regression Analysis of Seller Reports

	Report		
	(1)	(2)	(3)
Ethics Treatment	3.324	3.319	1.845*
	(2.673)	(2.717)	(1.047)
Asset Value			0.968***
			(0.015)
Age			0.169
			(0.151)
Female			0.805
			(1.162)
GPA			1.271
			(1.148)
Risk			0.206
			(0.289)
Observations	438	438	438
Adjusted R-squared	0.001	0.089	0.904
Round FE	No	Yes	Yes

Notes - This table summarizes the OLS regressions results of the amounts reported by the sellers on ethical treatment. The regressions are based on the sample with sellers only. The baseline is comprised of the amounts reported in the sessions with no ethical treatment. *Report* is the dollar amount the seller announces regarding the asset value. *Ethics Treatment* is an indicator variable that takes the values of 1 when sellers are provided with an ethical statement they have to sign; 0 otherwise. The control variables are comprised of individual seller characteristics (age, gender, GPA, and risk preference) as well as the true asset value in dollar amount. 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.

Table 5: Regression Analysis of Bids

	Bid		
	(1)	(2)	(3)
Ethics Treatment	9.531**	9.293**	8.380**
	(3.703)	(3.735)	(3.720)
Report			0.532***
			(0.035)
Age			-0.467
			(0.483)
Female			-4.621
			(3.563)
GPA			0.465
			(5.138)
Risk			0.818
			(0.816)
Observations	1,753	1,753	1,753
Adjusted R-squared	0.023	0.043	0.272
Round FE	No	Yes	Yes

Notes - This table summarizes the OLS regressions results of the amounts of bids by the bidders on ethical treatment. The regressions are based on the sample of bidders only. The baseline is comprised of the bids in the sessions with no ethical treatment. *Bid* is the dollar amount the bidder proposes. *Ethics Treatment* is an indicator variable that takes the value of 1 when sellers are provided with an ethical statement they have to sign (bidders observe the signed ethical statement, if any); 0 otherwise. The control variables are comprised of individual seller characteristics (age, gender, GPA, and risk preference) as well as the true asset value in dollar amount. Round FE indicates whether the estimates are within rounds. The standard errors are clustered at the bidder level, and reported in parentheses below the coefficients. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Table 6: Regression Analysis of Bidders Learning over the Rounds

	Prediction Error (Buyers)		
	(1)	(2)	(3)
Round x Ethics Treatment			0.174**
			-0.076
Round	-0.229***	-0.056	-0.227***
	(0.052)	(0.058)	(0.051)
Observations	997	756	1,753
Adjusted R-squared	0.123	0.081	0.119
Controls	Yes	Yes	Yes
Subject x Session FE	Yes	Yes	Yes
Sample	No Treatment	Treatment	Both Regimes

Notes - This table summarizes the OLS regressions results of bidder's prediction error over the rounds, with consideration of the differential in change in prediction error between mandatory and no ethical treatments. The specifications in Columns (1) and (2) follow equation 8 when considering each ethical treatment regime separately, while the specifications in column (3) follow equation 9 to capture the effects of mandatory ethical treatments on bidder's learning compared to the baseline of no ethical treatment. $Prediction\ Error$ is the absolute value of the difference between bidder's estimate of the asset value and the actual asset value. $Round$ is the number of rounds. *Ethics Treatment* is an indicator variable that takes the value of 1 if the seller is presented with an ethical statement to sign (mandatory treatment); 0 otherwise (when there are no ethical treatments). Controls is a vector comprised of individual bidders' characteristics (age, gender, GPA, and risk preference) as well as the reported amount. $Subject \times Session\ FE$ indicates that the estimates are within sessions for a given bidder. The coefficient on the stand-alone *Ethics* variable is omitted due to the $Subject \times Session$ fixed effect. The standard errors are clustered at the bidder level, and reported in parentheses below the coefficients. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Table 7: Regression Analysis: Relation between Ethics and Trust

	Report (Sellers)		*Bid (Buyers)*	
	(1)	(2)	(3)	(4)
Ethics Treatment	1.637	1.731	12.815***	3.361
	(1.246)	(1.743)	(4.345)	(5.945)
Observations	221	217	877	876
Adjusted R-squared	0.914	0.902	0.360	0.207
Controls	Yes	Yes	Yes	Yes
Round FE	Yes	Yes	Yes	Yes
Subsample	Low Trust	High Trust	Low Trust	High Trust

Notes - This table summarizes the OLS regressions cross-sectional results of the sellers' reports and bidders' bids on ethical treatment, when splitting each of the samples (bidders and sellers) 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. *Ethics Treatment* is an indicator variable that takes the value of 1 if the seller is presented with an ethics statement to sign (mandatory treatment); zero otherwise. Controls is a vector comprised of individual seller characteristics (age, gender, GPA, and risk preference) as well as the true asset value (for sellers sub-samples) or Value reported value (for bidders sub-samples). 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.

Table 8: Regression Analysis of Ethical Types

	Ehical type (Sellers)	
	(1) OLS	(2) Random effects
Ethics Treatment	0.010	-0.006
	(0.139)	(0.132)
Trend (Round #)	0.011	0.011
	(0.008)	(0.141)
Constant	0.681***	0.675***
	(0.168)	(0.161)
Observations	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. The estimation is conducted using OLS (column 1) and random effects (column 2). The variable *Ethics Statement* is an indicator variable equal to one when the seller must sign the ethical statement. 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.

Table 9: Counter-factual Analyses

	Avg price error	*Avg |price error|*	*Avg bias*	*Seller utility (e=1)*	*Seller utility (all e avg)*
Experiment (all treatments)	1.29	8.70	4.94	58.45	43.51
Nash ethics noise	0.20	6.66	8.57	50.21	46.45
Nash ethics noise (cost x 2)	0.08	3.71	4.59	48.56	50.20
Nash cash maximizers (e=1)	0.03	0.09	9.48	45.44	45.44
Second price auction	1.29	8.70	4.94	47.06	31.01

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 bidder'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 second 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. Second-price assumes that sellers bid their assessment and the second-highest price is set as the winning price.