NBER WORKING PAPER SERIES

# IS NO NEWS (PERCEIVED AS) BAD NEWS? AN EXPERIMENTAL INVESTIGATION OF INFORMATION DISCLOSURE

Ginger Zhe Jin
Michael Luca
Daniel Martin

Working Paper 21099
http://www.nber.org/papers/w21099

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2015

Patrick Rooney provided excellent research assistance. All errors are ours. The views expressed are those of the authors and do not necessarily represent those of the U.S. Federal Trade Commission, any individual Commissioner, or the National Bureau of Economic Research.

Is No News (Perceived as) Bad News? An Experimental Investigation of Information Disclosure
Ginger Zhe Jin, Michael Luca, and Daniel Martin
NBER Working Paper No. 21099
April 2015, Revised August 2016
JEL No. C9,D8,K2,L51

## ABSTRACT

This paper uses laboratory experiments to directly test a central prediction of disclosure theory: that market forces can lead businesses to voluntarily provide information about the quality of their products. This theoretical prediction is based on unraveling arguments, which require that consumers hold correct beliefs about non-disclosed information. Instead, we find that receivers are insufficiently skeptical about non-disclosed information, and as a consequence, senders do not always disclose their private information. However, when subjects are informed about non-disclosed information after each round, behavior slowly converges to full unraveling. This convergence appears to be driven by an asymmetric response in receiver actions after learning that they were profitably deceived. Despite the change in receiver behavior, stated beliefs about sender strategies remain insufficiently skeptical, which suggests that while direct and immediate feedback induces equilibrium behavior, it does not reduce strategic naïveté.

Ginger Zhe Jin
University of Maryland
Department of Economics
3115F Tydings Hall
College Park, MD 20742-7211
and NBER
jin@econ.umd.edu

Michael Luca
Harvard University
Soldiers Field
Boston, MA 02163
mluca@hbs.edu

Daniel Martin
Paris School of Economics
48 Boulevard Jourdan
75014 Paris
France
daniel@martinonline.org

# 1 Introduction

From the number of calories in a croissant to the fuel efficiency of a car, businesses routinely have private information about the quality of their products that potential customers would like to know. Businesses then face a decision – should they reveal or withhold this information? The goal of this paper is to shed light on the economics of voluntary disclosure through a series of laboratory experiments.

A central tenet of the economics of information is the idea that market forces can drive firms to voluntarily and completely disclose such information, as long as the information is verifiable and the costs of disclosure are small (Viscusi 1978, Grossman and Hart 1980, Grossman 1981, Milgrom 1981). The mechanism behind this idea is simple: consumers treat all non-disclosing companies the same, so the best businesses among those will have an incentive to separate themselves through disclosure. Applied iteratively, this logic produces unraveling in the quality of non-reporting firms, so that in equilibrium consumers correctly infer the very worst about information that is not disclosed. In other words, no news is bad news. The policy relevance of the unraveling result is clear – it highlights the potential for voluntary disclosure to solve asymmetric information problems across a variety of domains.

At the same time, the unraveling logic rests on restrictive assumptions around the ability of buyers to make inferences about a business's decision to withhold information. Unraveling can fail to occur when customers do not understand the implications of shrouded or missing information (Eyster and Rabin 2005, Gabaix and Laibson 2006, Mullainathan, Schwartzstein, and Shleifer 2008, Heidhues, Koszegi, and Murooka forthcoming).

Our aim is to investigate the unraveling predictions using lab experiments that are complex enough to capture the main strategic tensions of the theory yet simple enough for subjects to easily understand the structure of the game. In our experiments, there are two players: an information sender (e.g., the firm) and an information receiver (e.g., the consumer). The sender receives private information that perfectly identifies the true state (e.g., the firm's true quality level). The sender then makes a single decision: whether or not to disclose this information to the receiver. As a result, the sender cannot misrepresent the state.[1] By prohibiting dishonest reporting, we reproduce the assumptions underlying the unraveling prediction and mirror an important feature of many markets, such as those with truth-in-advertising laws.

[1] This is in contrast with existing experiments on strategic information transmission where senders can engage in "cheap talk" (Cai and Wang 2006, Wang, Spezio, and Camerer 2011).

After the sender decides whether or not to disclose her private information, the receiver must guess the state. If the sender has revealed the state, the receiver knows it with certainty. Otherwise, the receiver must infer the true state based on the sender's decision to withhold information and on the distribution of states, which is common knowledge. Reflecting many market transactions, the sender and receiver do not have aligned interests. The sender earns more when the receiver guesses that the state is higher (guesses and states are numeric values), and the receiver earns more when their guess is closer to the true state. With these payoffs, the logic of unraveling leads to the unique sequential equilibrium of this game: senders should always reveal their information (unless the state takes the lowest possible value, in which case they are indifferent between revealing and not), and receivers should correctly guess that the state takes the lowest possible value when senders do not reveal this information.

Based on the choices of 422 experimental subjects, we find a fundamental breakdown in the logic of unraveling: receivers are insufficiently skeptical about undisclosed information. That is, receivers underestimate the extent to which no news is bad news. This complements the growing field evidence on attention and inference in disclosure contexts. Brown, Hossain, and Morgan (2010) show that eBay customers are less elastic to shipping costs when they are listed separately from the base price. Brown, Camerer, and Lovallo (2012) find that firms with lower quality movies choose to engage in "cold openings" (i.e. they withhold movies from critics until the movie is released). Their data suggest that customers do not fully infer the fact that movies with cold openings tend to be worse.[2] Luca and Smith (2015) show that MBA programs with lower rankings are less likely to display them on their website. In these settings, there are other factors that can prevent disclosure and influence customer decisions, and our experiment allows us to strip away many of these factors and focus directly on the behavioral assumptions underlying unraveling. Our results also complement theoretical predictions of Eyster and Rabin (2005), Gabaix and Laibson (2006), and Heidhues, Koszegi, and Murooka (forthcoming).

To explore the barriers to unraveling, we have participants play our disclosure game 45 times, varying the extent of feedback we provide receivers during the game. In one group, receivers only learn the true state when it is revealed by the sender. In a second group, receivers are debriefed after each round and told the true state whether or not it was revealed by the sender. And in a third group, receivers are given aggregate information toward the end of the session about the overall likelihood of disclosure for each quality score. The results show that in group one, behavior never converges to the

[2] This result is firmly established in Brown, Camerer, and Lovallo (2013), where it is demonstrated how such data can be used to differentiate between equilibrium and non-equilibrium behavior in settings of verifiable disclosure (specifically related to the extent that naïveté limits unraveling).

full unraveling result. However, in group two, behavior slowly converges to the full unraveling result. Unlike round-by-round debriefs, the aggregate feedback given to group three does not lead to unraveling, highlighting the difficulty of correcting the biases of receivers.

We then explore the dynamics of disclosure across rounds and the cognitive barriers that prohibit subjects from playing equilibrium strategies. To this end, we elicit beliefs about aggregate sender strategies after all 45 rounds have been completed. In our experiment, feedback does not appear to impact beliefs about sender strategies, which suggests that the mechanism to convergence is not simply due to evolving beliefs about what others will play. At the same time, subjects that receive feedback do reduce their guesses significantly more after over-guessing than those that do not receive feedback. These facts suggest that subjects react to learning that they were profitably deceived, but not through a process of explicitly updating their subjective beliefs about the actions of others.

Overall, our results help to shed light on the economics of voluntary disclosure. Because immediate and direct feedback about non-disclosed information is often limited, there is reason to be skeptical that unraveling will occur in the field. For example, when restaurants choose not to disclose their hygiene ratings or calorie counts, it can be difficult, annoying, or time-consuming for consumers to assess this information about the restaurant immediately after having completed their meal. Moreover, this suggests that voluntary disclosure in markets with infrequent transactions may be less effective than in markets with frequent transactions for a given consumer, who then has more opportunity to learn.

In addition to offering a clean experimental design developed explicitly to investigate the predictions of disclosure theory, this paper provides three further contributions, which are discussed in more detail in the next section. First, we complement existing empirical studies from the field by providing direct evidence of naïveté in verifiable disclosure in a laboratory experiment.[3] Unlike existing studies from the field, the lab allows us to tightly control information and incentives, to minimize additional reasons for limited disclosure,[4] and to evaluate risk aversion and social preferences as alternative explanations.[5]

---

[3] Our relationship to existing verifiable disclosure experiments is proved in Section 2. Concurrent to our study, Hagenbach and Perez-Richet (2015) examine verifiable disclosure when payoffs to senders can vary.

[4] For a review of additional reasons for limited disclosure, see Dranove and Jin (2010). A partial summary is provided in Section 2.

[5] In addition, because we have subjects switch roles, we are able to show that many subjects who appear naïve in role of receiver nevertheless appear sophisticated in the role of sender. This suggests that receiver choices may be more "strategically complex" than sender choices. For instance, receivers need to undertake hypothetical thinking, which has been shown to be difficult in voting games (Esponda and Vespa 2015).

Second, we vary the feedback provided to subjects, and explore its role in producing convergence. The impact of feedback is especially policy relevant in this context, as policymakers have discussed various informational interventions related to disclosure. Our results suggest that absent immediate and direct feedback, consumers may not engage in the type of learning that induces full disclosure. Moreover, the impact of feedback depends on which form it takes – the type of aggregate feedback that is more common in the field does not improve inferences. This also sheds light on contexts in which voluntary disclosure may occur and suggests unraveling is unlikely to happen in settings with limited feedback about quality (such as the number of calories in a sandwich) or with fewer transactions (such as house or car purchases).

Third, we elicit beliefs about sender strategies, which yields direct evidence that feedback may not correct such beliefs. This suggests that naïveté in disclosure might persist even with experience in the marketplace and the behavioral bias from strategic naïveté could be alleviated by payoff-based learning.

The rest of the paper is organized as follows. Section 2 describes three related literatures and our contributions to each. Section 3 lays out the disclosure game and its equilibrium features. Section 4 describes our experimental design. Section 5 reports the extent to which the unraveling prediction occurs in our experiment with and without feedback. Section 6 explores the role of feedback and beliefs in convergence to theoretical predictions. A brief discussion of the results is offered in Section 7.

## 2 Related Literature

Our paper draws on and contributes to three literatures: the literature on voluntary disclosure, the literature on communication experiments, and the literature on beliefs and learning in games.

### 2.1 Voluntary Disclosure

Voluntary disclosure is appealing from a policy perspective because it can improve consumer welfare even without mandatory disclosure policies, which are often opposed by industry groups and challenging to implement and enforce. The classic unraveling result suggests that the same benefits as mandatory disclosure can be achieved simply by ensuring that disclosed information is verifiable and the related costs are low. This has inspired a number of measures, including standardized information displays, certification agencies, and truth-in-advertising laws.

In practice, voluntary disclosure is observed in many industries, but is far from complete.[6] As summarized in Dranove and Jin (2010), this incompleteness has motivated two strands of theories to account for why unraveling does not occur. One strand emphasizes external factors such as disclosure cost and consumer knowledge before disclosure, while the other strand focuses on a seller's strategic incentives. As an example of the latter, sellers may choose not to obtain data on product quality in order to avoid future demand for disclosure (Matthews and Postlewaite 1985).

Other examples of strategic incentives include product differentiation (Board 2009) and countersignaling with multiple quality dimensions (Feltovich, Harbaugh, and To 2002). The seller's strategic incentive can also be dynamic: one may refrain from disclosure even if he has favorable information at hand, as he fears that today's disclosure may make it harder to explain non-disclosure in the future when the information turns out non-favorable (Grubb 2011). In another example of dynamic incentives, a pharmaceutical firm may prefer to be silent about the potential health risks of its products because of litigation risk, but this may crowd out positive disclosures (Marinovic and Varas 2015).

In this paper, we use lab experiments to exclude these additional reasons for limited disclosure and therefore create a test environment closer to the original, classical theory. The strategic incentives for disclosure that we study are also present in the subsequent literature, so our results can potentially inform the wider literature as well. For instance, persistent naïveté about non-disclosure could be combined with any of the additional forces given above to produce new predictions for verifiable disclosure.

Our work also draws on the literature in behavioral economics, which has posited that if buyers are naïve about the quality of non-disclosed information, then sellers may not disclose all of their private information. Eyster and Rabin (2005) consider this possibility in the context of their "cursed equilibrium" concept, and Gabaix and Laibson (2006) and Heidhues, Koszegi, and Murooka (2012) consider it in the context of shrouded attributes. Mullainathan, Schwartzstein, and Shleifer (2008) present a model of coarse thinking that highlights how informational spillovers from one environment to another can make non-disclosure more persuasive than it would otherwise be. Likewise, in the accounting literature, Hirshleifer and Teoh (2003) consider the impact of naïveté on financial disclosures. In their model, receivers can be naïve to non-disclosed information, but they can also be inattentive to disclosed information.

We add to these literatures in several ways. First, we provide evidence of naïveté in voluntary disclosure through a controlled laboratory experiment where beliefs are

[6] See Mathios 2000, Jin 2005, Fung et al. 2007, and Luca and Smith 2015 for specific examples.

elicited. Second, we show that without immediate feedback on non-disclosed information, such naïveté can persist over time and is not easily eliminated even if subjects are provided information about aggregate disclosure behavior. Third, we show that subjects who appear strategically sophisticated in the role of sender can also appear strategically naïve in the role of receiver.

## 2.2 Communication Experiments

Our design borrows many features from the cheap talk experiments of Cai and Wang (2006) and Wang, Spezio, and Camerer (2010). For instance, we follow both of these experiments in describing the sender's type using "secret" numbers and in starting messages to the receiver with "The number I received is." In addition, our type space and payoffs are similar to those found in Wang, Spezio, and Camerer (2010).

The key difference in our experimental design is that the sender's messages must be truthful. Hence, our experiment tests models of verifiable disclosure, rather than cheap talk. There are only a limited number of experiments that include verifiable disclosure, and there are important design differences between these experiments and ours. Three of these papers (Forsythe, Isaac, and Palfrey 1989, King and Wallin 1991, and Dickhaut, Ledyard, Mukherji, and Sapra 2003) are focused on disclosure in asset markets (as in Milgrom and Roberts 1986). These experiments feature a sender (the asset seller) who decides whether to disclose the asset's quality to receivers who compete with each other through an auction mechanism. Forsythe, Isaac, and Palfrey (1989) find "unravelling of both the prices paid for blind-bid items and the quality levels of these items." King and Wallin (1991) and Dickhaut, Ledyard, Mukherji, and Sapra (2003) complement these findings by also showing what happens when there is the possibility that senders may not be informed about the asset's quality. The latter goes beyond the former by considering both partially informed senders and partially informative messages. These experiments represent a valuable test of disclosure in asset markets, but the use of auctions (particularly first-price auctions) introduces room for other biases to drive disclosure decisions.

In addition, Forsythe, Lundholm, and Rietz (1999) compare disclosure to cheap talk in reducing adverse selection, and they find that reports converge to the unraveling predictions. Their verifiable disclosure treatment differs from our experiments in that receivers have a more complicated choice (what price to ask for the product), and senders can choose not to take that price.

Concurrent to but separate from our study are two new papers that use experiments to study verifiable disclosure. Benndorf, Kübler, and Normann (2015) study a disclosure game in a labor market setting where multiple senders compete through the use of disclosure, but unlike in our experiments, the receiver is a computer that uses an automated strategy, so there is no room for inference problems. Hagenbach and Perez-Richet (2015) investigate a verifiable disclosure game where sender payoffs are not necessarily monotonic in the state space.

In two experiments that study lying aversion, senders have three options: tell the truth, lie, or not disclose. Non-disclosure takes the form of vague messages in the case of Serra-Garcia, van Damme, and Potters (2011) and silence in the case of Sanchez-Pages and Vorsatz (2009), so the latter is closer to our experiment. However, unlike in our experiments, in Sanchez-Pages and Vorsatz (2009) non-disclosure carries a cost. Even with this cost, some senders choose not to disclose. Serra-Garcia, van Damme, and Potters (2011) find that intermediate senders sometimes use vague messages, which receivers do not make correct inferences about. Agranov and Schotter (2012) also study the use of vague language but focus on the vagueness possible with human language.

Relative to this literature, we believe that our experiments present a simpler and more direct test of classic verifiable disclosure theory. In addition, to the best of our knowledge, our experiment is the first to vary the feedback provided to subjects, to elicit beliefs about sender strategies, or to provide information about aggregate sender behavior.

### 2.3 Beliefs, Learning, and Equilibrium

Central to any strategic interaction is the set of beliefs that people hold about each other. This has given rise to two related questions within the experimental literature. First, do people hold correct beliefs about how other people play and do they best respond to these beliefs? Second, to what extent do people dynamically learn about markets and strategic interactions?

While economists typically infer beliefs from actions, stated beliefs can provide further evidence on both the belief formation process and the ways in which people react to their own stated subjective beliefs. For example, Costa-Gomes and Weizsacker (2008) find that subjects often do not best respond to stated beliefs about sender strategies, while Rey Biel (2007) finds much higher rates of best responding to these beliefs in simpler games.

Building on this literature, we ask subjects to state their beliefs about sender strategies. This allows us to compare stated beliefs across treatment arms and also to see whether receivers are best responding to their own beliefs. We find a strong, positive, and statistically significant correlation between implied beliefs about the underlying state (based on beliefs about sender strategies) and receiver guesses of the underlying state, which suggests that their actions incorporate beliefs about sender strategies.

The fact that our results do not converge to the static equilibrium gives rise to questions about the dynamics of belief formation. The experimental literature on learning in games has focused broadly on two types of learning. Models of belief-based learning assume that people form beliefs about how others will play in strategic contexts, and choose actions based on those beliefs. In contrast, models of payoff-based learning, in particular reinforcement-based learning, do not require people to explicitly form expectations – and instead assume that people learn by simply by reacting to observed payoffs, which in the long run will lead toward actions that yield higher payoffs. Feltovich (2000) offers a comparison of these two types of learning models in games of asymmetric information. In an experimental "horse race" he finds that both outperform Nash equilibrium in explaining behavior and that neither is better than the other across all metrics. Camerer and Ho (1999) offer an approach called Experience-Weighted Attraction (EWA) that captures aspects of both types of models.

Building on this, we present evidence against belief-based learning when we show that beliefs about sender strategies are not statistically different when subjects are given feedback and when they are not given feedback. However, subjects who receive feedback reduce their guesses after over-guessing more than those that do not receive feedback, which is suggestive of payoff-based learning. At the margin, beliefs do not seem to update in response to new information, but strategies do seem to react to payoff information.

## 3 The Disclosure Game

The one-shot disclosure game we study involves two agents: an information sender and an information receiver. At the beginning of the game, nature determines the state $s$ (which can be interpreted as the sender's type) by taking a draw from a probability distribution $F$ with full support over a finite state space $S$, which is a subset of the real numbers. The sender knows the realized state, but ex ante, the receiver knows only the distribution of possible states.

The sender has two possible actions, and the receiver is aware that these are the only two actions available to the sender. The sender can either report the state to the receiver or make no report. This report must be truthful and cannot be vague. Thus, the set of actions

*M* available to a sender of type $s$ is just *M(s)={s,null}*.[7]

Regardless of whether or not they receive a report from the sender, the receiver takes an action $a$ from a finite space $A$, which is also a subset of the real numbers and contains S. We interpret this action as guessing the type of the sender.

The true state and the receiver's action determine the payoffs for the two parties. The sender's utility is given by a function $U_S(a)$, which is concave, monotonically increasing in the receiver's action, and independent of the state. The receiver's utility is given by a function $U_R(a,s)$, which is concave in the receiver's action and reaches its peak when $a$ is equal to $s$. In other words, the receiver benefits the most from selecting an action that is as close as possible to the true state, while the sender benefits the most when the receiver's action is as high as possible. These utility functions produce a strong conflict of interest when the state is low.

When the set of receiver actions $A$ is sufficiently rich, the techniques found in Milgrom (1981) can easily be adapted to show that in every sequential equilibrium of this disclosure game, the sender always reports the state (unless it is the minimum element in S), and if there is no report, the receiver takes the action that is the minimum element in S. In other words, the sender always reports his or her type (unless it is the worst possible type), and the receiver always guesses the sender is the worst possible type if they do not report. When the realized state is the minimum element in S, the sender is indifferent between reporting or not, so any mixture over these actions is consistent with equilibrium.

There are other Bayesian Nash equilibria of this game, but they require strategies to contain behavior off of the equilibrium path that is not sequentially rational.[8] Importantly, we do not observe behavior consistent with any other Bayesian Nash equilibrium when we implement this game experimentally.

## 4 Experimental Design

In our experiments, subjects completed 45 rounds and then, depending on the session, one of six possible additional tasks. Subjects were told at the beginning of the experiment

---

[7] In the model of Milgrom (1981), senders are allowed to report a range of states, but we consider a simpler message space in order to reduce the strategic complexity of the game, which could add cofounding factors.

[8] For instance, when the set of receiver actions is sufficiently rich there can exist a Nash equilibrium in which sellers never report and receivers take the action that is as close to the average realization of the state space as possible. This is supported by a receiver strategy in which the action equal to the minimum element of the state space is taken if the sender does report.

that they would complete an additional task, but were given no details about the task. The appendix contains the full set of instructions given before the start of the experiment. Instructions for each additional task were presented to subjects on the computer screen just before the start of the task.

At the end of each session, subjects were paid, privately and in cash, their show-up fee plus any additional earnings from the experiment. Over the course of the experiment, subjects had the opportunity to accumulate or lose "Experimental Currency Units" (or ECUs). At the end of the experiment, each subject's ECU balance was rounded up to the nearest non-negative multiple of 200 and converted into U.S. dollars at a rate of 200 to 1.

### 4.1 In Each Round

In each round, subjects were randomly matched into pairs. To reduce reputational effects, subjects were matched anonymously and were told that it was very unlikely they would be paired with the same subject in consecutive rounds. For a session size of 14, the actual likelihood of being paired with the same subject in consecutive rounds is 0.6%.

In each round and for each pairing, one subject was randomly assigned to be the sender, and the other subject was assigned to be the receiver. Each was equally likely to be assigned either role. As a result, the likelihood of a subject experiencing both roles by round 5 is 93.75%. We used alternating roles to ensure that receivers understood that senders could not misreport the state. To reduce framing effects, the sender was referred to as the "S Player," and the receiver was referred to as the "R Player."

For each pair, the computer drew a whole number from 1 to 5, called the "secret" number. Thus, the state space was $S=\{1,2,3,4,5\}$. Each of these numbers was equally likely to be drawn, and both senders and receivers were made aware of this probability distribution over the state space.

Each sender was shown the secret number for their pairing and then made their decision while the receivers waited. Senders were given the option to either "report" or "skip", with no time limit on their decision.

After all senders made their decisions, the receivers' screens became active. If a sender decided to report their secret number, the receiver they were paired with was shown this message: "The number I received is," followed by the actual secret number. If a sender decided instead to skip any reporting, the area for messages on the receiver's screen was left blank. Subjects were told that these were the only two actions available to senders, so if the area for messages on the receiver's screen was left blank, it was because the sender chose not report the secret number.

Below the area for messages, receivers were asked to guess the secret number, and these guesses could be any half unit between 1 and 5. Thus, the set of actions is A={1,1.5,2,2.5,3,3.5,4,4.5,5}. The action space of receivers was made sufficiently rich that the sequential equilibrium involves full unraveling. There was also no time limit for receiver decisions.

Subjects were paid for every round, and receiver payoffs in each round were $ECU_R = 110 - 20|S - A|^{1.4}$, where S is the secret number and A is the receiver's guess. These payoffs are such that a risk neutral receiver would guess closest to their expected value of the secret number. The sender payoffs in each round were $ECU_s = 110 - 20|5 - A|^{1.4}$. These payoffs are independent of the secret number and monotonically increasing with receiver actions, because guesses could not be higher than 5. These payoffs are similar to the quadratic specification found in Crawford and Sobel (1982) when there is a large bias towards higher actions. Because we use just a small number of states and actions, the payoffs could be shown in a table, so that subjects did not need to know or interpret these functional forms.

With these payoff functions, there was a clear misalignment of interests between senders and receivers. Receiver payoffs were higher when their guesses were closer to the secret number, and sender payoffs were higher when the receiver made higher guesses. Subjects were told about these two broad features of sender and receiver payoffs.

## 4.2 After Each Round

Our primary treatment variation was in the information provided as feedback to subjects after each round. In our "no feedback" treatment, subjects were given no information after completing each round. After all receivers had made their decisions, subjects proceeded to a screen that required them to click "OK" to start the next round. After all subjects had pressed this button, the next round began.

In contrast, in our "feedback" treatment, subjects were told four pieces of information after each round: the actual secret number, whether the sender reported the secret number, the receiver's guess of the secret number, and their own payoff. After all subjects pressed the "OK" button on the screen containing this feedback, the next round began.

To reduce social considerations, subjects in the feedback treatment were not told the payoffs for the other player in their pairing, though it could be deduced using the payoff table. To reduce wealth effects and portfolio effects, they were also not told their cumulative payoffs earned so far in the experiment. In addition, subjects only received feedback about their pairing, not all pairings in the session.

## 4.3 Belief Elicitation and Additional Tasks

After completing 45 rounds, subjects completed one of six additional tasks. Before undertaking this additional task, some subjects were asked to guess the rate at which senders reported each secret number in the initial 45 rounds of the experiment. The purpose of this question was to assess whether subject beliefs about sender strategies were correct. These guesses were not incentivized, but in a recent paper, Trautmann and Kuilen (2015) show that such "introspective" elicitation can yield accurate beliefs.

In the first of six additional tasks, which we call the "Risk" task, subjects completed the well-known measure of risk aversion introduced by Holt and Laury (2002). The aim of this task was to see whether sender and receiver choices were related to the risk preferences of subjects.

For this measure, subjects make 10 choices between a "safer" lottery (payments of \$2.00 or \$1.60) and "riskier" lottery (payments of \$3.85 or \$0.10) in which the probability of the high payment was the same within each choice, but varied across choices. A risk-neutral decision maker would choose the lottery with a 40% chance of \$2 over the lottery with a 40% chance of \$3.85, but the lottery with a 50% chance of \$3.85 over the lottery with a 50% chance of \$2. The switching point in this "multiple price list" can be viewed as a reflection of the risk preferences of each subject. This task was incentivized by randomly selecting one of their 10 choices, realizing the chosen lottery, and adding any earnings to the show-up fee and earnings from the first 45 rounds.

We call the second additional task the "Other" task. In this task, subjects played once more in the role of sender and once more in the role of receiver, but in both cases, they played against a computer instead of a human (and were told this was the case). This computer played a strategy designed to mimic the past decision of another player. This type of task is designed to keep the strategic decisions the same as in previous choices, but to remove the playoff implications for others. By comparing these choices with previous choices, we can determine whether sender and receiver choices were impacted by any social preferences. Niederle and Vesterlund (2007) use a similar approach to separate preferences for competition from social preferences.

For this task, when in the role of receiver, subjects were told that the computer sender would not report the secret number and that the secret number would be from a random past round in which the secret number was not reported. When in the role of sender, subjects were told that if they reported the secret number, the computer receiver would guess that number, and if they did not report, the computer receiver would repeat the guess of a receiver from a random past round where the secret number was not reported. To get as much information as possible from the sender decision, we used the “strategy” method in which senders made a decision for each possible secret number before seeing the actual secret number. The payoffs from this task were added to the ECU earned in the first 45 rounds.

In the third task, which we call the “Self” task, subjects played once more in the role of sender and once more in the role of receiver, and in both cases, they also played against a computer instead of a human. However, this time the computer played a strategy designed to mimic the past decisions of that same subject. This type of task is designed to assess whether subjects can best respond to accurate beliefs, under the assumption that they form accurate beliefs about their own strategies. A similar approach was used by Ivanov, Levin, and Niederle (2010) in examining the role of beliefs in the Winner’s Curse.

For this task, when in the role of receiver, subjects were told that the computer sender would not report the secret number and that the secret number would be from a random past round in which they did not report the secret number. When in the role of sender, subjects were told that if they reported the secret number, the computer receiver would guess that number, and if they did not report, the computer receiver would repeat their own guess from a random past round where the secret number was not reported. In this task, we also used the “strategy” method for sender decisions. The payoffs from this task were added to the ECU earned in the first 45 rounds.

In the fourth task, which we call the “Computer” task, subjects played 5 additional rounds in the role of receiver against a computer sender. In this task, subjects were told that the S player (computer) would report the secret number if that would “maximize their earnings given the guesses of all other participants (besides yourself) in the proceeding round.” In practice, this meant that the computer reported the secret number if it was above the average guess for all other subjects in the previous round who did not receive a report. The payoffs from this task were added to the ECU earned in the first 45 rounds. The aim of this task was to assess whether any failures of unraveling in the first 45 rounds were due solely to the fact that receivers believe senders were potentially non-optimizing or poorly informed humans, which may be a good assumption for small firms, but not necessarily large firms.

With the last two additional tasks, we tested two possible informational interventions. Subjects were shown some information about the play of all subjects in the first 45 rounds, then completed the same choices as in the Other task, and then played 5 more rounds just as in the first 45 rounds. The payoffs from this task were added to the ECU earned in the first 45 rounds.

In the fifth additional task, which we call the "Average Reports" task, the information that subjects were shown was the average *reported* secret number from all subjects in that session from the first 45 rounds. However, because the number of rounds in which the secret number was reported was not provided, there was not enough information for subjects to infer anything about the average non-reported secret number.

We call the final additional task the "Disclosure Rates" task because subjects were shown the number of times that each secret number was reported and not reported for all subjects from the first 45 rounds. This provided enough information to determine both the average reported secret number and the average non-reported secret number.

## 4.4 Main Analysis Sample

Our study was conducted in the Center for Experimental Social Science (CESS) laboratory at New York University (NYU) and the Computer Lab for Experimental Research (CLER) facility at the Harvard Business School (HBS). In both laboratories, subjects are separated with dividers, and each subject is provided with a personal computer terminal. Our experiment was run using the z-Tree software package (Fischbacher 2007).

Across schools, we observed 422 subjects complete a total of 18,976 decisions, which corresponds to 9,488 pairings in 45 rounds.[9] Over 34 sessions, the median session size was 12 subjects and the mode session size was 14. At both schools we used a show-up fee of $5, and on average subjects earned $25.25 at NYU and $25.41 at HBS. Across schools, the minimum payment was $14 and the maximum payment was $30.85.

At HBS, 120 subjects were assigned to sessions where no feedback was provided and 90 were assigned to feedback sessions. All 210 subjects were asked for their beliefs about sender strategies and completed the Disclosure Rates additional task. At NYU, all 212 subjects were assigned to the no feedback sessions. Of these 212 subjects, none were

[9] Due to a problem encountered by the program during the final round of one session, receiver guesses are missing for 7 pairings. We do not include these observations in the analysis sample.

asked for their beliefs about sender strategies, 38 completed the Risk additional task, 26 the Other task, 38 the Self task, 42 the Computer task, 30 the Average Reports task, and 38 the Disclosure Rates task.[10] These additional tasks are used for robustness checks, so are excluded from the main analysis sample.

Table 1 presents summary statistics for all 18,976 decisions. For 21.3% of those decisions, feedback was provided after each round, which corresponds to our feedback treatment. For 64.8% of decisions, the secret number was reported in the pairing, and the unconditional average guess of the secret number in the pairing was 3.145. The number of decisions in both schools was roughly equal: 50.2% at NYU and 49.8% at HBS. The NYU and HBS subjects differ in terms of demographics: 63.3% of the HBS decisions were made by undergraduates,[11] 49% were made by males, 86.2% by native English speakers (1.4% declined to report), and 10% by subjects who reported having a friend present in the room during the session. In comparison, the NYU subjects are more likely to self report as undergraduates, female, and non-native English speakers.

# 5 Main Results: Unraveling in the Lab

This section presents our main results. When feedback is not provided, or is provided at aggregated levels, we find that behavior does not converge towards equilibrium predictions – even after many rounds. This is because receivers do not sufficiently infer that no news is bad news. However, when detailed feedback is provided after each round, receiver guesses and sender disclosure rates move towards equilibrium predictions.

## 5.1 Failure of Unraveling

Overall, our results suggest that unraveling often fails to occur, even when full unraveling is the prediction. This is driven by the fact that receivers are insufficiently skeptical about the undisclosed information. However, frequent and full feedback can push behavior toward the equilibrium prediction.

Table 2 summarizes sender and receiver actions. As shown in Panel A, senders are more likely to report higher draws. In both no feedback and feedback treatments, the average reporting rate is above 80% when the draw is equal to the population average (a secret number of 3), and jumps to about 95% when the draw is 4 or 5. For lower draws,

[10] These variations reflect our changing access to experimental labs over time. We will discuss the similarity and difference between NYU and HBS sessions later in the text.

[11] At both schools, subjects were required to be 18 years old. At CESS, subjects were restricted to be students, and in the feedback sessions at CLER, subjects were required to be less than 26 years old.

however, the average reporting rate drops to 42.69% for 2 and 8.61% for 1. In total, senders choose silence 35.2% of the time, which is a sharp contrast to the unraveling prediction.[12]

Between no feedback and feedback, the sender reporting rate differentiates most for lower draws. When the draw is 2, 50.25% of senders will report under feedback, which is significantly higher than the 40.69% who report under no feedback. Similarly, when the draw is 1, 11.16% of senders report under feedback, but only 7.89% report under no feedback.

Because all feedback sessions were run in the HBS lab, we further decompose no feedback observations into HBS and NYU separately. Compared with HBS senders, NYU senders are less likely to report when the draw is low (1, 2, 3) but more likely to report when the draw is above average (4, 5). As shown later, this is primarily driven by the demographic difference between NYU and HBS. For a cleaner comparison, the last column of Table 2 tests whether the sender reporting rate is statistically different between no feedback and feedback observations for HBS subjects only. This comparison is even starker than in the full sample: HBS senders in the feedback regime are more likely to report all intermediate draws (2, 3, 4), and the biggest difference occurs when the draw is 2 (50.25% in feedback versus 42.57% in no feedback). For extreme draws (1 and 5), the reporting rate is higher in feedback observations, but no longer statistically significant from no feedback.

Turning to receivers, Panel B presents the average receiver guess by feedback treatments, conditional on whether the sender reports 1, 2, 3, 4, 5, or blank. Because senders are not allowed to misreport, one may expect receivers to guess exactly the reported number if the sender discloses it. This expectation is largely confirmed when the reported number is 3 or 4, but with some deviation when the reported number is close to either extreme. In particular, receivers tend to over-guess at the low extreme (1, 2) and under-guess at the high extreme (5). While a boundary effect may be one explanation, we believe a more plausible explanation has to do with the social preferences of the receivers. Because senders tend to be silent when the draw is close to the low extreme (1, 2), receivers may want to reward reporting out of reciprocity, even if over-guessing reduces receiver payoff. In contrast, when the reported number is at the high end (5), some receivers may punish the sender's good luck by under-guessing. These social preferences seem stronger in no feedback than in feedback, and more likely to occur in HBS than in NYU, which also seems to be due to demographic differences.

---

[12] In the unraveling equilibrium, only senders with the lowest draw may choose non-disclosure, which is expected to be around 20% of draws in our experiment.

The most striking pattern occurs when the sender chooses non-disclosure. In this situation, receivers guess 2.16 with no feedback and 1.90 with feedback. Relative to the truth, receivers over-guess by 27.8% with no feedback and 13.1% with feedback.

The difference in guesses between no feedback and feedback is always statistically significant, no matter whether the comparison is run in the whole sample or conditional on HBS only. Furthermore, if we combine Panel A with Panel B, it is clear that receivers over-guess when the secret number is not reported, even under feedback.

### 5.2 Round-by-Round Dynamics

Next, we analyze the dynamics of behavior across 45 rounds and with varying levels of feedback. Overall, we find that actions converge toward equilibrium when subjects receive round-by-round feedback, which is consistent with a large number of papers that have demonstrated movement towards equilibrium across rounds in experiments. However, there are a number of subtleties in the dynamics of behavior, including a larger reduction in mistakes by receivers than by senders, even though subjects play both roles.

As shown in Figure 1, the average sender reporting rate increases from 56-58% in the first five rounds to 67-74% in the last five rounds. This increase is steady and similar for no feedback and feedback. A closer look suggests that the growth is most apparent for the draws of 2 and 3, with faster growth in feedback than in no feedback when the draw is 2 (Figure 1B).

Figure 1C plots the average receiver guess round by round, conditional on sender non-disclosure. Under no feedback, the average receiver guess drops sharply in the first five rounds, but stays in the range of 2 to 2.2 afterwards. This contrasts with the steady decline of the average draw of blank reports, which reaches as low as 1.54 in the last five rounds. The persistent gap between receiver guesses and real draws suggests that, under no feedback, the game does not converge to the unraveling equilibrium even after 45 rounds of repeat play. In comparison, the average receiver guess of blank reports drops steadily under feedback, to 1.43 in the last five rounds. This is getting close to the average draw of blank report in these rounds (1.33), suggesting a tendency to slowly converge toward the unraveling equilibrium under feedback.

These dynamic patterns are further confirmed in regressions, where we control for subject demographics or subject fixed effects. As shown in Table 3, senders learn round by round, especially when the draw is 2 in feedback. Facing a blank report from senders, receivers learn as well. The speed of receiver learning nearly doubles when the treatment is feedback (versus no feedback). According to the last column, regression estimates suggest that the average receiver guess of blank reports drops sharply from 2.39 in round 1 to 1.90 in round 45 for no feedback and to 1.34 in round 45 for feedback. Again, this suggests slow convergence towards the unraveling equilibrium, but for feedback only.

The disclosure equilibrium discussed in Section 3 is based on the concept of Nash equilibrium: each player takes the opponent's strategy as given; each player chooses the best action in response to the opponent's strategy; and each player, if facing incomplete information, forms a rational expectation that confirms the true, hidden information on average. To better understand the degree of disequilibrium in our experiment, we calculate the best action a player should adopt in response to what the opponent has played in the past round(s). Table 4 summarizes the absolute mistake, as measured by the distance between the computed best response and the player's real action. For robustness in how many historical rounds players use in learning, we compute the responding period by both one round (Panel A) and five rounds (Panel B). In theory, the absolute mistake should diminish to zero as players learn and converge to the equilibrium strategies.

Consistent with Figure 1C and Table 3, the absolute mistakes of receivers do drop significantly from 0.8426 for HBS with no feedback to 0.6860 for HBS with feedback (in response to senders in the past five rounds). Even if we pool NYU and HBS for no feedback, the receiver mistake remains significantly smaller in feedback than in no feedback. Similar patterns arise when we define the responding period as one round instead of five rounds. This robustness suggests that feedback plays an important role in reducing receiver mistake. In comparison, the reduction of sender mistake is only significant from no feedback to feedback when we focus on HBS subjects and define the responding period as five rounds.

Table 5 examines whether the absolute mistake decreases over time, while controlling for either subject demographics or subject fixed effects. For receivers, the round-by-round reduction in absolute mistakes is more apparent in feedback than in no feedback. For senders, mistakes decline every period in both feedback regimes, although the mistakes relative to the best response in the past five rounds seems to reduce more in feedback. Overall, Table 4 and Table 5 confirm what we have seen in the raw data: both senders and receivers tend to learn over time, but feedback speeds up receiver learning. Put another way, receiver over-guessing persists more in no feedback than in feedback, although senders continue to increase reporting rate across all 45 rounds. Because receivers and senders are exactly the same people by definition, this suggests that the receiver's problem is more difficult and somehow feedback helps to overcome this difficulty.

### 5.3 Robustness Checks

To reinforce our interpretation of the results as a failure of inferences, we show that our results are not driven by three competing explanations – the risk preferences of receivers, social preferences, and the size of the state space. In this section, we provide evidence against these explanations.

First, to investigate the role that risk aversion plays in receiver choices, we use a standard measurement tool from experimental economics for assessing risk preferences. As described in the previous section, 38 NYU subjects from our no feedback treatment completed the Risk task after the initial 45 rounds. When a subject has more than one switch point in the Holt-Laury multiple price list, then risk preferences are hard to ascertain, but just 3 subjects had multiple switch points.

For the 35 subjects that had consistent switch points, 5 had a switch point that is consistent with risk neutrality. Another 3 subjects had switch points consistent with being risk loving, and the rest of subjects were consistent with being risk averse. There was a fair bit of variation in switch points: 5 subjects switched from the safe lottery to the risky lottery when there was a 50% chance of the high payment, 8 switched when there was a 60% chance, 7 when there was a 70% chance, and 5 when there was an 80% chance.

We used an OLS regression of receiver guess onto switch point to look for evidence of positive relationship between risk aversion and the size of guesses. Controlling for the number of rounds that a receiver had spent as a sender or receiver up to that point and for subject fixed effects, the coefficient on switch point is indeed positive, but is small (0.031) and not significant ($p=0.261$). In this analysis, the unit of observation is at the subject level.

Second, we explore the role of social preferences. There are many ways that social preferences can influence play in our experiment, but many, such as altruism on the part of senders or rewards for disclosure from receivers, would push behavior towards equilibrium. One that could push behavior away from equilibrium is if receivers guess higher than they would otherwise to reduce the imbalance in payoffs between senders and receivers. Because of the concavity of the payoff function, when receivers make very low guesses, sender payoffs are very low. In many standard social preference models, agents lose utility when they experience guilt over making much higher payoffs than their opponent. Such models would predict that receivers would make higher guesses, even when the secret number is not reported.

For evidence of this, we examine 26 NYU subjects from the no feedback treatment who completed the Other additional task. As mentioned previously, these subjects guessed the secret number from an earlier round, but without playoff implications for the sender. If social preferences were a leading explanation for higher guesses, we would expect a decrease in guesses in this task. Instead, the average guess increased by 0.206, which is not statistically significant at a 10% level (two-sided t-test, p=0.1866). Once again, the unit of observation is at the subject level.

This increase in guesses after social motives are minimized provides suggestive evidence of a punishment motive towards those who do not disclose. Instead of force pushing away from equilibrium, the social element in choice (which could be due to social preferences) appears to be pushing behavior towards equilibrium. In other words, to the extent that social preference exists, it appears to produce a punishment motive. This does not explain why receivers tend to over-guess when the secret number is not reported, nor does it explain why feedback tends to reduce over-guessing.

As demonstrated before, receivers do respond to past plays that they have experienced earlier in the session. Based on this, one may argue that many subjects may be fully aware of the unraveling equilibrium in theory but they do not follow the equilibrium receiver strategy because they believe their opponents are not sophisticated enough to follow the equilibrium sender strategy. To provide some robustness along this dimension, we use computerized senders to approximate sophisticated opponents. In our "Computer" task, receivers switch from playing a human sender to playing a computer sender, and as described previously, the computer plays optimally given the choices of all other receivers in the past round. This should reveal the extent to which choices are shaped by playing human senders.[13]

For the 34 NYU subjects in our no feedback treatment who completed this task, we find that receiver guesses decrease slightly, but not enough to make full disclosure a best response for all computer senders. The difference between the average guess when playing against human senders and the average guess when playing against the computer sender is statistically significant at the 5% level (two-sided t-test, p=0.0021), but is just 0.196. Again, the unit of observation is at the subject level.

Because computer senders considered the behavior of *other* receivers when deciding whether to report to a receiver, the likelihood of reporting was not the same for all computer senders, even within the same session. When the secret number was 2,

[13] This task also changes the social considerations, but we have found these to be negligible in our experiment using the Other task.

computer senders disclosed the number around 56.8% of the time, close to the rate of human senders.

Third, one may argue that the limited space of five states may make it difficult to conclude for or against the unraveling prediction, because there is only one state (2) that is strictly below average but above the worst draw. To get a sense for how the size of the state space might impact our findings, we ran an experiment with a space of 10 states – S={1,2,3,4,5,6,7,8,9,10} – which is twice as large as in the original experiment. Here again we allow receivers to guess half-unit intervals, so the action space is A={1,1.5,2,2.5,…,9,9.5,10}.

To keep payoffs in a similar range to the original experiment, the distance from the ideal action is divided in half in the payoff functions, so that receiver payoffs are $ECU_R = 110 - 20|(S - A)/2|^{1.4}$ and sender payoffs are $ECU_S = 110 - 20|(10 - A)/2|^{1.4}$. As a result, the payoffs for senders and receivers when the receiver guesses 4 and the state is 2 is the same in the new experiment as when the receiver guesses 2 and the state is 1 in the original experiment.

Aside from increasing the set of secret numbers and changing the payoff table, the experimental design and instructions are the same as in the original experiment with no feedback. We conducted this experiment in the CLER facility at HBS. 84 subjects completed the new experiment, and the median and mode session size were both 14. Once again, we limited the subject pool to be 25 years old or younger.

Table 6 provides the summary of player actions in both the no feedback sessions of the main analysis sample for both NYU and HBS and in all sessions of the 10-state experiment. Panel A shows the average reporting rate by secret number in both experiments. As in the primary study with 5 secret numbers, the reporting rate increases monotonically with the secret number in the experiment with 10 secret numbers. The reporting rate for a secret number of 3 in the new experiment is 41.21%, which is comparable to and not statistically different from the reporting rate for a secret number of 2 in the primary study of 40.69%. In addition, the reporting rate in the new experiment for secret numbers of 5, 7, and 9 are similar and not statistically different from the reporting rates for 3, 4, and 5 in the primary study.

A secret number of 3 in new experiment and a secret number of 2 in primary study are also comparable in the sense that a risk neutral sender is close to indifferent between reporting and not reporting at a secret number of 3 in the new experiment. The average guess for a non-reported secret number is 3.419 with a 95% confidence interval of 3.252 to 3.584. As in the primary study, the average guess is above the average actual non-reported secret number in the new experiment. In the new experiment, the average non-report secret number is 2.616, which is 0.803 below the average guess. We also find a similar pattern in over and under-guessing for reported secret numbers. In both experiments, reporting lower secret numbers is on average rewarded with a higher guess, while reporting the highest secret number is on average punished with a lower guess.

In short, we conclude that failure of unraveling – as measured by incomplete reporting by senders and over-guessing of blank reports by receivers – is robust to enlargement of state space.

# 6 Beliefs, Feedback, and Behavior

In this section, we explore the mechanism driving the feedback effect. Essentially, there are two broad reasons that feedback might improve decision-making. First, it could correct beliefs about sender actions through what is known as belief-based learning. Second, it could be that subjects are changing their behavior in response to realized payoffs, but without explicitly forming strategic beliefs. See Section 2.3 for a brief review of the related literature.

To investigate these possibilities, we take three steps. First, we elicit subjects' stated beliefs about how often senders of each type report their private information. Second, we compare the responses of receivers to different kinds of news that feedback can contain. Third, we provide subjects with aggregate information about sender behavior, both to directly correct mistaken beliefs and to mimic the type of information available to consumers in some markets.

## 6.1 Stated Beliefs about Sender Strategies

As described in Section 4, after all 45 rounds were completed and before any additional tasks were undertaken, we asked all 210 subjects at HBS to guess the percentage of senders over all 45 rounds who reported each secret number. The frequency of their responses is given in Figure 2.

Larger bubbles correspond to a larger number of guesses in an interval of 5 percentage points, and the solid line corresponds to the average guess. While there is clear heterogeneity in the stated beliefs of subjects, the bulk of guesses follow the average rate. One exception is that there are a large number of guesses above the average guess for a secret number of 3. However, the two dashed lines, which represent the actual reporting rate, do pass through this region. The upper dashed line represents the reporting rate for the last 20 rounds, and the lower dashed line represents the average reporting rate for all 45 rounds. Though subjects were asked to guess the reporting rate for all 45 rounds, it seems plausible that their guesses would be closer to the actual reporting rate for the more recent rounds. However, for all secret numbers, the average guesses are indeed closer to the average reporting rate for all 45 rounds.

## 6.2 Few Differences in Stated Beliefs between Treatments

Of the 210 subjects that were asked for their beliefs about sender strategies, 90 received round-by-round feedback and 120 received no feedback between rounds. Because subjects were asked for their belief data after all 45 rounds were complete, if feedback helps subjects to learn sender strategies, we should see a difference between the stated beliefs of these subjects. Table 7 provides just such a comparison. As the top panel of the table shows, the only significant difference between treatments in reporting rate occurs in guesses for a secret number of 3 (for a two-sided t-test). Although this difference is statistically significant, the magnitude of the difference is less than 10 percentage points. On top of this, reporting rates are actually higher in the feedback treatment, so even without learning, a strategically aware agent should guess higher in the feedback treatment.

In addition, if we look at what these reported beliefs imply about non-disclosed secret numbers, there is no significant difference between treatments. If subjects correctly use Bayes' rule, then combining their beliefs about sender strategies with the probability of each secret number implies a belief about non-reported secret numbers. For example, if a receiver believes that the sender reporting rate is 20% for a secret number of 1, 40% for 2, 80% for 3, and 100% for 4 and 5, then this implies that the secret number conditional on non-reporting is 1.625 because [(100%-20%)*1+(100%-40%)*2+(100%-80%)*3]/[(100%-20%)+(100%-40%)+(100%-80%)]=1.625. We refer to this implied belief about non-reported secret numbers as the "implied guess." The average implied guess is strikingly similar between no feedback (1.952) and feedback (1.888), and the difference is not significant by any statistical standard (p=0.4072).

## 6.3 Stated Beliefs, Observed Actions, and Actual Guesses

Regardless of whether feedback is given, another channel for belief-based learning is what receivers observe in sender actions – specifically how often each secret number is disclosed. The information contained in sender actions is provided in the second panel of Table 7. The panel shows the percentage of each report, including blank ones, observed by receivers on average for each treatment. Using Bayes' Rule, these observations imply something about what receivers should guess, and this number is also given. Note that what receivers should guess based on these observations is significantly different between treatments at the 10% level.

The third panel shows what subjects actually guessed on average in each treatment and across all rounds, for the first 25 rounds, and for the final 20 rounds. If this kind of observational learning impacts the strategic beliefs of receivers, then we should see a strong correlation between the implied guess based on sender strategies and what people should guess given what they observe. The final panel shows that this is not the case. The correlation is just 0.103 for the no feedback treatment and 0.207 in the feedback treatment. In addition, the correlations between what people actually guess and what they should guess based on these observations are not statistically significant.

Finally, if beliefs about actual reporting rates explains why receivers are not suitably skeptical of non-disclosed information on average, then we would expect to see a strong correlation between what reported beliefs about sender strategies imply about the secret number and what people actually guess for the secret number. These correlations are significant at the 1% level, which suggests both that beliefs about sender strategies is an important factor in the strategic naïveté of receivers and that receivers incorporate their beliefs about sender strategies.

### 6.4 Asymmetric Responses to Feedback

Though we do not have evidence of feedback impacting beliefs about sender strategies, it is possible that feedback allows receivers to learn to make better guesses by showing them which guesses have higher payoffs. To look for evidence of this, we examine whether changes from one guess of a non-disclosed secret number to the next guess of a non-disclosed secret number are related to the types of mistakes made and whether feedback was received. Table 8 shows the results of several regression specifications based on this objective. Regressions 1 through 3 do not include subject fixed effects, but regressions 4 through 6 do include them. Regressions 1, 2, 4, and 5 just consider the direction of the mistake in the prior guess, while regressions 3 and 6 consider both the magnitude and the direction.

Across all regression specifications, we find strong evidence that subjects who were informed that they guessed too high decreased their guesses the next time they had an opportunity to do so, and that this effect is both statistically significant at the 1% level and substantial. If subjects have access to other sources of learning, we would expect this effect to exist for both treatments, and it does. However, the effect is stronger for subjects receiving feedback for all specifications except one, where the point estimate indicates that the effect is stronger, but not statistically significant.

On the other hand, the impact of feedback is not symmetric. There is not a statistically significant relationship between receiving feedback and under-guessing for 5 of the 6 specifications, and in the one specification where it is significant, it is just significant at the 10% level. Moreover, the point estimates of the effect sizes are much smaller than for over-guessing.

### 6.5 Little Effect from Aggregate Reporting

The final question we address is whether providing information about aggregate sender strategies impacts behavior. As shown previously, we find evidence that reported beliefs about sender strategies are not skeptical enough and are strongly correlated with actual guesses, so information about sender strategies could potentially improve guesses and disclosure rates.

To test this, we examine the choices of subjects who completed the Average Reports and Disclosure Rates tasks. As mentioned previously, after 45 rounds the subjects are shown the corresponding aggregate information. We then have subjects play 5 more rounds, and we compare the reporting rates and guesses in these rounds to the reporting rates and guesses made in the last 5 rounds of the first 45 rounds. We do this for the 30 NYU subjects in the no feedback treatment who completed the Average Reports task, the 38 NYU subjects in the no feedback treatment who completed the Disclosure Rates task, the 120 HBS subjects in the no feedback treatment who completed the Disclosure Rates task, and the 90 HBS subjects in the feedback treatment who completed the Disclosure Rates task. These results are provided in Table 9.

The top panel of the table compares the reporting rates for the 5 rounds just before and just after the information is provided. For senders in the feedback treatment, the reporting rates do not significantly differ for a two-sided t-test, nor does the average secret number when senders make no disclosure. The same is true if we look at both the Average Reports subjects and the Disclosure Rates subjects in the no feedback treatment.

The bottom panel compares the guesses made when senders did and did not report their secret number. Once again, there are no significant differences for either of the informational interventions or for either of the treatments. If anything, the Disclosure Rates treatment causes guesses to rise on average.

# 7 Discussion

Our findings shed light on a fundamental problem preventing full information unraveling. In our experiments, receivers are not sufficiently skeptical about undisclosed information. The extent to which these mistakes persist depends on the extent and type of feedback provided to receivers. When receivers get no feedback or aggregated feedback, this optimism persists through the full 45 rounds of the experiment. In these contexts, information senders profit by limiting disclosure.

However, round-by-round feedback about mistakes does debias receivers and lead behavior to converge to the sequential equilibrium prediction. Our results also provide insights on the mechanism underlying learning about disclosure. Round-by-round feedback does not fix the wrong belief about sender strategy, but it does foster a learning process that helps receivers to make lower and lower guesses – consistent with the idea of payoff-based learning.

Our results also shed light on the factors that may limit voluntary disclosure in the field, and the situations in which we might expect voluntary disclosure to be an effective policy. These findings suggest that market forces are insufficient to close the information gap between sellers and buyers, unless buyers receive fast and precise feedback about mistakes after each transaction.

For the products that naturally offer such feedback – say cereals that taste crunchy and t-shirts that hold color fast – voluntary disclosure may converge to the unraveling results after each buyer purchases the product many times. However, for product attributes with less immediate feedback – such as the fat content of salad dressing and the cleanliness of a restaurant kitchen – voluntary disclosure may not converge to the unraveling results. In these situations, mandatory disclosure may be necessary if the policy goal is complete disclosure.

# 8 References


Agranov, M., & Schotter, A. (2012). Ignorance is bliss: An experimental study of the use of ambiguity and vagueness in the coordination games with asymmetric payoffs. *American Economic Journal: Microeconomics*, *4*(2), 77–103.

Battigalli, P., & Guaitoli, D. (1997). Conjectural equilibria and rationalizability in a game with incomplete information. In *Decisions, games and markets* (pp. 97–124). Boston: Kluwer American Publishers.

Benndorf, V., Kübler, D., & Normann, H. T. (2015). Privacy concerns, voluntary disclosure of information, and unraveling: An experiment. *European Economic Review*, *75*, 43–59.

Board, O. (2009). Competition and disclosure. *Journal of Industrial Economics*, *57*(1), 197–213.

Brown, A. L., Camerer, C. F., & Lovallo, D. (2012). To review or not to review? Limited strategic thinking at the movie box office. *American Economic Journal: Microeconomics*, *4*(2), 1–26.

Brown, A. L., Camerer, C. F., & Lovallo, D. (2013). Estimating structural models of equilibrium and cognitive hierarchy thinking in the field: The case of withheld movie critic reviews. *Management Science*, *59*(3), 733–747.

Brown, J., Hossain, T., & Morgan, J. (2010). Shrouded attributes and information suppression: Evidence from the field. *The Quarterly Journal of Economics*, *125*(2), 859–876.

Cai, H., & Wang, J. T. Y. (2006). Overcommunication in strategic information transmission games. *Games and Economic Behavior*, *56*(1), 7–36.

Camerer, C., & Hua Ho, T. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, *67*(4), 827–874.

Costa-Gomes, M. A., & Weizsäcker, G. (2008). Stated beliefs and play in normal-form games. *The Review of Economic Studies*, *75*(3), 729–762.

Crawford, V., & Sobel, J. (1982). Strategic information transmission. *Econometrica*, *50*(6), 1431–1451.

Dekel, E., Fudenberg, D., & Levine, D. K. (2004). Learning to play Bayesian games. *Games and Economic Behavior*, *46*(2), 282–303.

Dickhaut, J., Ledyard, M., Mukherji, A., & Sapra, H. (2003). Information management and valuation: an experimental investigation. *Games and Economic Behavior*, *44*(1), 26–53.

Dranove, D., & Jin, G. Z. (2010). Quality disclosure and certification: Theory and evidence. *Journal of Economic Literature*, *48*(4), 935–963.

Esponda, I. (2008). Behavioral equilibrium in economies with adverse selection. *American Economic Review*, *98*(4), 1269–91.

Esponda, I., & Vespa, E. (2015). Endogenous sample selection and partial naiveté: A laboratory study. Mimeo.

Eyster, E., & Rabin, M. (2005). Cursed equilibrium. *Econometrica*, *73*(5), 1623–1672.

Feltovich, N. (2000). Reinforcemet-based vs. belief-based learning models in experimental asymmetrc-information games. *Econometrica*, *68*(3), 605–641.

Feltovich, N., Harbaugh, R., & To, T. (2002). Too cool for school? Signalling and countersignalling. *RAND Journal of Economics*, *33*(4), 630–649.

Fischbacher, U. (2007). z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics*, *10*(2), 171–178.

Forsythe, R., Isaac, R. M., & Palfrey, T. R. (1989). Theories and tests of "blind bidding" in sealed-bid auctions. *RAND Journal of Economics*, *20*(2), 214–238.

Forsythe, R., Lundholm, R., & Rietz, T. (1999). Cheap talk, fraud, and adverse selection in financial markets: Some experimental evidence. *Review of Financial Studies*, *12*(3), 481–518.

Fung, A., Graham, M., & Weil, D. (2007). *Full Disclosure: The Perils and Promise of Transparency*. Cambridge and New York: Cambridge University Press.

Gabaix, X., & Laibson, D. I. (2006). Shrounded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics*, *121*(2), 505–540.

Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *Journal of Law and Economics, 24*(3), 461–483.

Grossman, S. J., & Hart, O. D. (1980). Disclosure laws and takeover bids. *Journal of Finance*, *35*(2), 323–334.

Grubb, M. (2011). Developing a reputation for reticence. *Journal of Economics & Management Strategy*, *20*(1), 225–268.

Hagenbach, J., & Perez-Richet, E. (2015) communication with evidence in the lab. Mimeo.

Heidhues, P., Koszegi, B., & Murooka, T. (forthcoming). Inferior products and profitable deception. *Review of Economic Studies*.

Hirshleifer, D., & Teoh, S. H. (2003). Limited attention, information disclosure, and financial reporting. *Journal of Accounting and Economics*, *36*(1), 337–386.

Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, *92*(5), 1644–1655.

Ivanov, A., Levin, D., & Niederle, M. (2010). Can relaxation of beliefs rationalize the winner's curse?: An experimental study. *Econometrica*, 78(4), 1435–1452.

Jin, G. Z. (2005). Competition and disclosure incentives: An empirical study of HMOs. *RAND Journal of Economics*, 93–112.

Kessler, J. B., & Roth, A. E. (2012). Organ allocation policy and the decision to donate. *American Economic Review, 102*(5), 2018–47.

King, R. R., & Wallin, D. E. (1991). Voluntary disclosures when seller's level of information is unknown. *Journal of Accounting Research*, *29*(1), 96–108.

Luca, M., & Smith, J. (2015). Strategic disclosure: The case of business school rankings. *Journal of Economic Behavior & Organization*, 112, 17–25.

Marinovic, I., & Varas, F. (2015). No news is good news: Voluntary disclosure in the face of litigation. *Stanford University Graduate School of Business Research Paper No. 13–19.*

Mathios, A. D. (2000). The impact of mandatory disclosure laws on product choices: An analysis of the salad dressing market. *Journal of Law and Economics*, *43*(2), 651–677.

Matthews, S., & Postlewaite, A. (1985). Quality testing and disclosure. *RAND Journal of Economics*, *16*(3), 328–340.

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

Milgrom, P., & Roberts, J. (1986). Relying on the information of interested parties. *RAND Journal of Economics*, *17*(1), 18–32.

Mullainathan, S., Schwartzstein, J., & Shleifer, A. (2008). Coarse thinking and persuasion. *The Quarterly Journal of Economics*, *123*(2), 577-619.

Niederle, M., & Vesterlund, L. (2007). Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics*, *122*(3), 1067–1101.

Rey-Biel, P. (2007). Equilibrium play and best response to (stated) beliefs in constant sum games (No. 676.07). *Unitat de Fonaments de l'Anàlisi Econòmica (UAB) and Institut d'Anàlisi Econòmica (CSIC) Working Papers.*

Sánchez-Pagés, S., & Vorsatz, M. (2009). Enjoy the silence: An experiment on truth-telling. *Experimental Economics*, *12*(2), 220–241.

Serra-Garcia, M., van Damme, E., & Potters, J. (2011). Hiding an inconvenient truth: Lies and vagueness. *Games and Economic Behavior*, *73*(1), 244–261.

Trautmann, S. T., & Kuilen, G. (2015). Belief elicitation: A horse race among truth serums. *The Economic Journal*, *125*(589), 2116–2135.

Viscusi, W. K. (1978). A note on “lemons” markets with quality certification. *Bell Journal of Economics*, *9*(1), 277–79.

Wang, J. T. Y., Spezio, M., & Camerer, C. F. (2010). Pinocchio's pupil: Using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *American Economic Review*, *100*(3), 984–1007.

# 9 Appendix: Experimental Instructions

**Welcome**

You are about to participate in an experiment on decision-making, and you will be paid for your participation in cash, privately at the end of the experiment. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance.

Please silence and put away your cellular phones now.

The entire session will take place through your computer terminal. Please do not talk or in any way communicate with other participants during the session.

We will start with a brief instruction period. During the instruction period you will be given a description of the main features of the experiment and will be shown how to use the computers. If you have any questions during this period, raise your hand and your question will be answered so everyone can hear.

**Instructions**

The experiment you are participating in consists of 45 rounds. At the end of the final round, you will complete an additional task, be asked to fill out a questionnaire, and then will be paid the total amount you have accumulated during the course of the session (in addition to the $5 show up fee). Everybody will be paid in private. You are under no obligation to tell others how much you earned.

The currency used during these 45 rounds is what we call "Experimental Currency Units" (ECU). For your final payment, your earnings during these 45 rounds will be converted into US dollars at the ratio of 200:1 (200 ECU=$1). They will then be rounded up to the nearest (non-negative) dollar amount.

In the first round, you will be matched with one other person, and you are equally likely to be matched with any other person in the room. You will not know whom you are matched with, nor will the person who is matched with you. One of you will be assigned to be S Player and the other to be the R Player for that round. You are equally likely to be assigned to either role. In the second round, you will once again be randomly matched with one other person (most likely with a different person than in the first round) and randomly assigned a role, and this will be repeated until 45 rounds are complete.

In each round and for every pair, the computer program will generate a secret number that is randomly drawn from the set {1,2,3,4,5}. The computer will then send the secret number to the S Player. After receiving this number, the S Player will choose whether or not to report the secret number to the R Player. If the S Player chooses to report the number, the R Player will receive this message from the S Player: "The number I received is" followed by the actual secret number. Otherwise, the R Player will receive no message.

After seeing the message or not, the R Player will guess the value of the secret number. The earnings of both players depend on the value of the secret number and the R Player's guess.

The specific earnings are shown in the table below, which is displayed again before the S Player and R Player make their choices. In each cell of the table, the payoff for the S Player is on the left, and the payoff for the R Player is on the right. As you can see from the table, the S Player earns more when the R Player makes a higher guess, and the R Player earns more when their guess is closer to the secret number.

| PAYOFFS: S , R | R's guess: 1 | R's guess: 1.5 | R's guess: 2 | R's guess: 2.5 | R's guess: 3 | R's guess: 3.5 | R's guess: 4 | R's guess: 4.5 | R's guess: 5 |
|---|---|---|---|---|---|---|---|---|---|
| Secret number: 1 | -29 , 110 | -6 , 102 | 17 , 90 | 38 , 75 | 57 , 57 | 75 , 38 | 90 , 17 | 102 , -6 | 110 , -29 |
| Secret number: 2 | -29 , 90 | -6 , 102 | 17 , 110 | 38 , 102 | 57 , 90 | 75 , 75 | 90 , 57 | 102 , 38 | 110 , 17 |
| Secret number: 3 | -29 , 57 | -6 , 75 | 17 , 90 | 38 , 102 | 57 , 110 | 75 , 102 | 90 , 90 | 102 , 75 | 110 , 57 |
| Secret number: 4 | -29 , 17 | -6 , 38 | 17 , 57 | 38 , 75 | 57 , 90 | 75 , 102 | 90 , 110 | 102 , 102 | 110 , 90 |
| Secret number: 5 | -29 , -29 | -6 , -6 | 17 , 17 | 38 , 38 | 57 , 57 | 75 , 75 | 90 , 90 | 102 , 102 | 110 , 110 |

Figure 1A: Sender disclosure rates by round

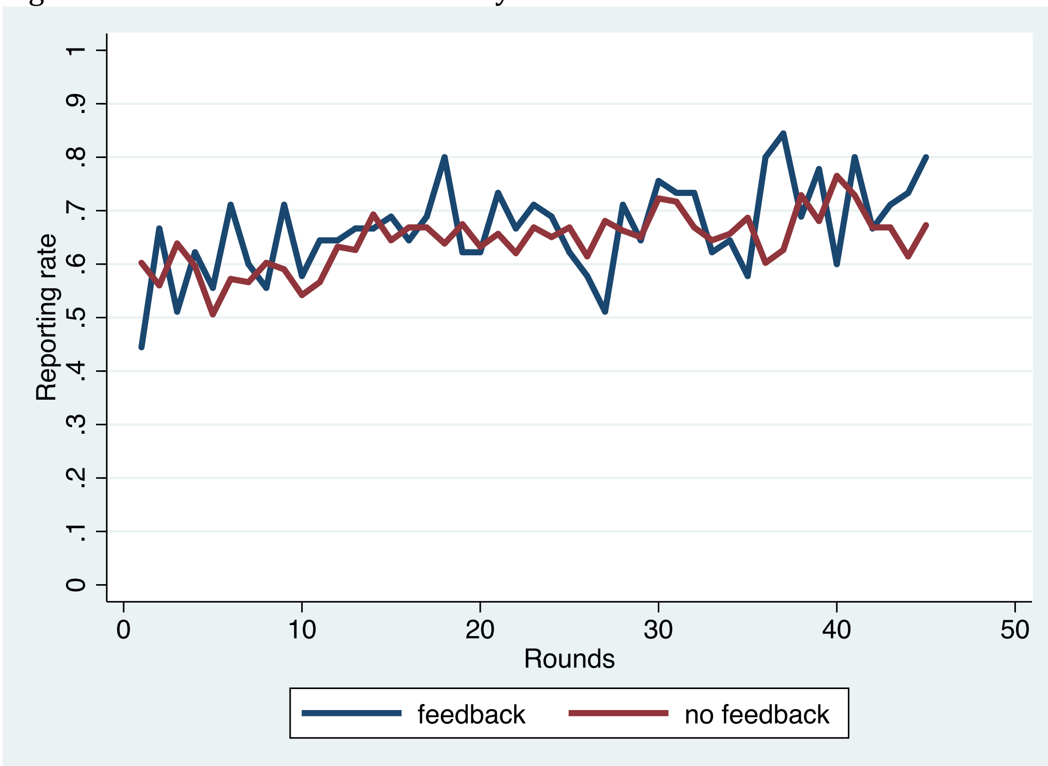


Figure 1B: Sender disclosure rate by round for draw of 1, 2, or 3

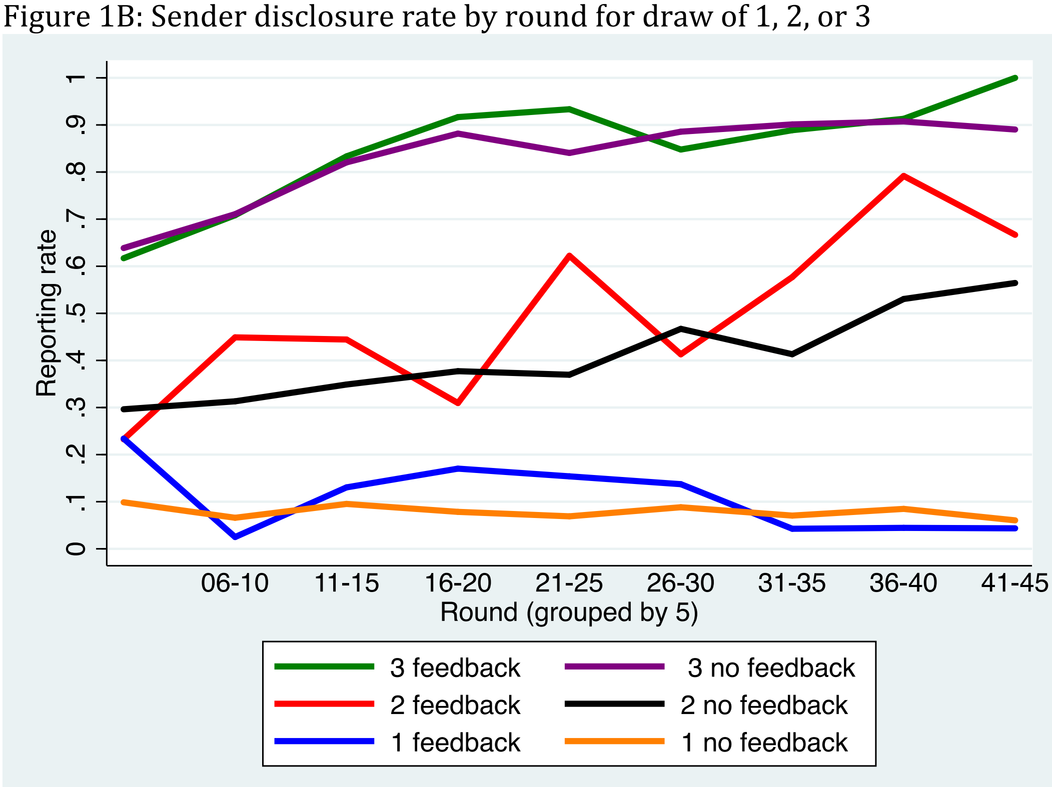

Figure 1C: Average receiver guess of non-disclosed numbers by round

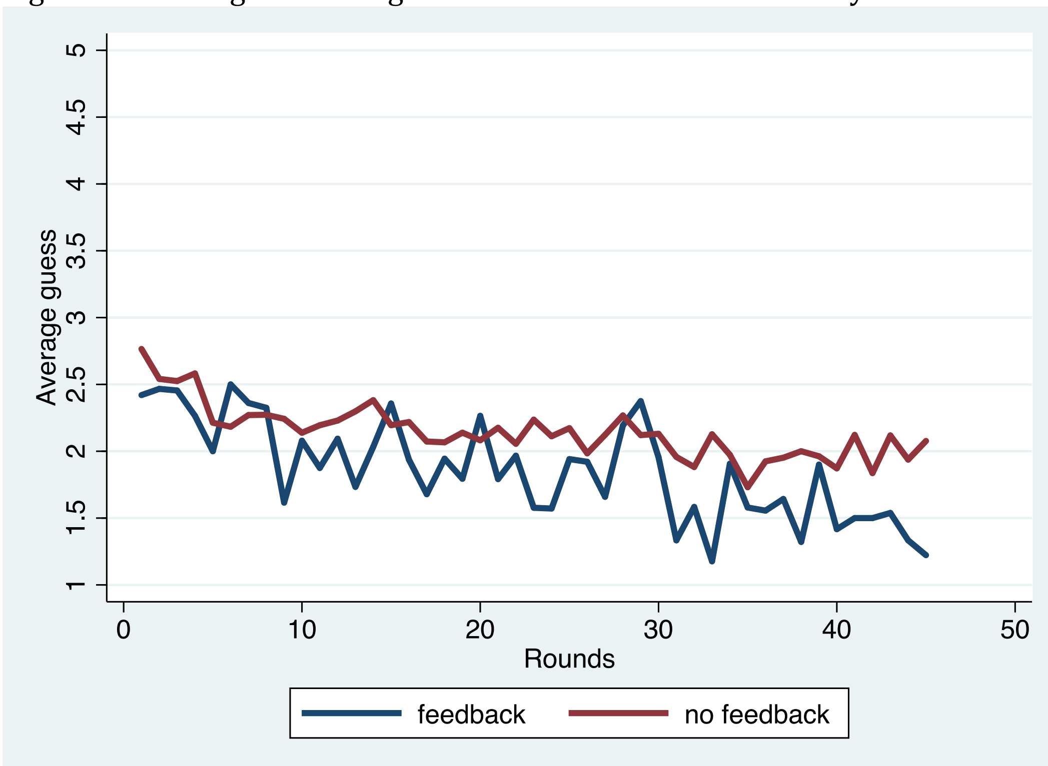


Figure 2: Receiver guess of sender strategies vs. actual disclosure rates

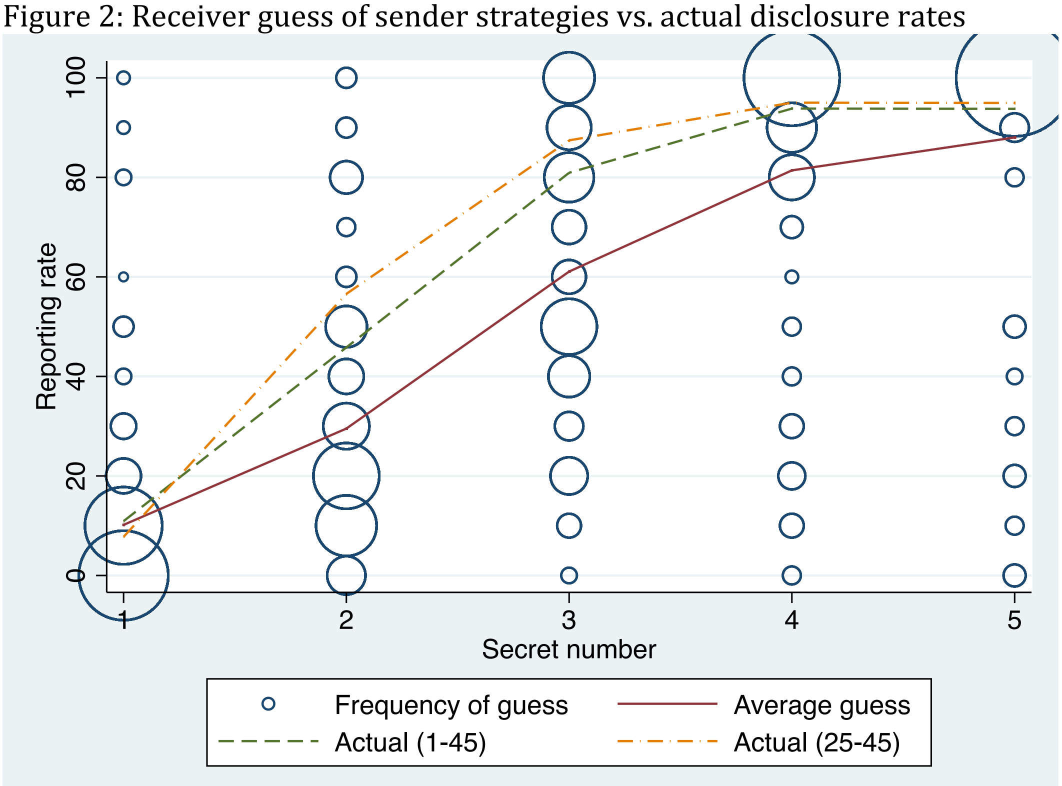

**Table 1: Sample summary**

Observation is per subject and per round
Conditional on rounds<=45, draws=1,2,3,4,5

| Variable | Obs. | Mean | Std. | Min | Max |
|---|---|---|---|---|---|
| Session ID | 18,976 | 16.800 | 9.745 | 1 | 34 |
| # of subjects in the session | 18,976 | 13.013 | 2.459 | 6 | 18 |
| Feedback treatment | 18,976 | 0.213 | 0.410 | 0 | 1 |
| Report | 18,976 | 0.648 | 0.478 | 0 | 1 |
| Guess | 18,976 | 3.145 | 1.298 | 1 | 5 |
| Period | 18,976 | 22.984 | 12.979 | 1 | 45 |
| Dummy=1 if subject is from NYU | 18,976 | 0.502 | 0.500 | 0 | 1 |
| Undergraduate (year<=4) | 18,976 | 0.746 | 0.435 | 0 | 1 |
| Male | 18,976 | 0.429 | 0.495 | 0 | 1 |
| Native English speaker | 18,976 | 0.780 | 0.414 | 0 | 1 |
| Dummy=1 if the subject does not Report being a native speaker | 18,976 | 0.007 | 0.084 | 0 | 1 |
| Having a friend in the same session | 18,976 | 0.097 | 0.296 | 0 | 1 |
| HBS only: | | | | | |
| Undergraduate (year<=4) | 9,450 | 0.633 | 0.482 | 0 | 1 |
| Male | 9,450 | 0.490 | 0.500 | 0 | 1 |
| Native English speaker | 9,450 | 0.862 | 0.345 | 0 | 1 |
| Dummy=1 if the subject does not report being a native speaker | 9,450 | 0.014 | 0.119 | 0 | 1 |
| Having a friend in the same session | 9,450 | 0.100 | 0.300 | 0 | 1 |
| NYU only: | | | | | |
| Undergraduate (year<=4) | 9,526 | 0.858 | 0.349 | 0 | 1 |
| Male | 9,526 | 0.368 | 0.482 | 0 | 1 |
| Native English speaker | 9,526 | 0.698 | 0.459 | 0 | 1 |
| Dummy=1 if the subject does not report being a native speaker | 9,526 | 0.000 | 0.000 | 0 | 0 |
| Having a friend in the same session | 9,526 | 0.094 | 0.292 | 0 | 1 |

* If a subject does not report "Native English", we code it as zero.

**Table 2: Summary of player actions**

| | | Panel A: Sender Report? (0/1) | | | | | |
|---|---|---|---|---|---|---|---|
| Sender's draw | Frequency | No Feedback (HBS or NYU) | Feedback (HBS) | p-value (=) | No Feedback (HBS) | No Feedback (NYU) | p-value (=, HBS only) |
| 1 | 1916 | 7.89% | 11.16% | 0.0346** | 10.74% | 6.15% | 0.8326 |
| 2 | 1942 | 40.69% | 50.25% | 0.0005*** | 42.57% | 39.63% | 0.0185** |
| 3 | 1879 | 83.23% | 84.75% | 0.4677 | 77.91% | 86.01% | 0.0092*** |
| 4 | 1882 | 95.86% | 95.56% | 0.7923 | 92.59% | 97.71% | 0.0650* |
| 5 | 1869 | 95.60% | 94.94% | 0.5703 | 92.87% | 97.18% | 0.1911 |
| Obs. | 9488 | 7463 | 2025 | | 2700 | 5985 | |
| Average draw if not reported | 3337 | 1.69 | 1.68 | 0.7129 | 1.80 | 1.63 | 0.0148** |

| | | Panel B: Receiver Guess | | | | | |
|---|---|---|---|---|---|---|---|
| Sender's message | Frequency | No Feedback (HBS or NYU) | Feedback (HBS) | p-value (=) | No Feedback (HBS) | No Feedback (NYU) | p-value (=, HBS only) |
| 1 | 165 | 1.38 | 1.3 | 0.5777 | 1.53 | 1.22 | 0.2183 |
| 2 | 829 | 2.15 | 2.16 | 0.9035 | 2.24 | 2.1 | 0.1298 |
| 3 | 1570 | 3.05 | 3.07 | 0.2103 | 3.06 | 3.04 | 0.7323 |
| 4 | 1803 | 4.01 | 4.02 | 0.763 | 4.01 | 4.01 | 0.7625 |
| 5 | 1784 | 4.91 | 4.97 | 0.0222** | 4.83 | 4.96 | 0.0001*** |
| Blank | 3337 | 2.16 | 1.90 | 0.0000*** | 2.28 | 2.08 | 0.0000*** |
| Obs. if reported | 6151 | 4801 | 1350 | | 1686 | 3115 | |
| Obs. if not reported | 3337 | 2662 | 675 | | 1014 | 1648 | |
| Obs. total | 9488 | 7463 | 2025 | | 2700 | 4763 | |

Note: *** p<0.01; ** p<0.05; * p<0.1.

**Table 3: Regressions on convergence**

| Variables | Sender reports? (0/1) | | | | Receiver guess of non-report | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Constant | -0.00613 | 0.0200 | 0.0241** | 0.0474*** | 2.770*** | 2.394*** |
| | (0.0152) | (0.0151) | (0.00937) | (0.00941) | (0.0610) | (0.0233) |
| Draw=2 | 0.343*** | 0.220*** | 0.345*** | 0.228*** | | |
| | (0.0129) | (0.0263) | (0.0119) | (0.0244) | | |
| Draw=3 | 0.749*** | 0.749*** | 0.745*** | 0.745*** | | |
| | (0.0106) | (0.0106) | (0.0102) | (0.0102) | | |
| Draw=4 | 0.873*** | 0.872*** | 0.881*** | 0.880*** | | |
| | (0.00806) | (0.00800) | (0.00827) | (0.00820) | | |
| Draw=5 | 0.871*** | 0.871*** | 0.868*** | 0.868*** | | |
| | (0.00815) | (0.00809) | (0.00855) | (0.00850) | | |
| Round | 0.00268*** | 0.00175*** | 0.00242*** | 0.00156*** | -0.0119*** | -0.0110*** |
| | (0.000299) | (0.000285) | (0.000286) | (0.000284) | (0.00134) | (0.000933) |
| Feedback treatment | 0.0148 | 0.0110 | | | -0.122 | |
| | (0.0196) | (0.0199) | | | (0.0801) | |
| Round * feedback | 0.000701 | 0.000110 | 0.000975 | 0.000266 | -0.00987*** | -0.0124*** |
| | (0.000655) | (0.000630) | (0.000632) | (0.000623) | (0.00283) | (0.00236) |
| Round * draw=2 | | 0.00459*** | | 0.00425*** | | |
| | | (0.000998) | | (0.000900) | | |
| Feedback * draw=2 | | 0.00568 | | 0.00254 | | |
| | | (0.0578) | | (0.0534) | | |
| Round * feedback * draw=2 | | 0.00343 | | 0.00401** | | |
| | | (0.00217) | | (0.00200) | | |
| Undergrad | 0.0505*** | 0.0491*** | | | -0.447*** | |
| | (0.00911) | (0.00908) | | | (0.0428) | |
| Native English | -0.0323*** | -0.0315*** | | | 0.161*** | |
| | (0.00867) | (0.00863) | | | (0.0383) | |
| Have friend in the session | -0.0193* | -0.0191* | | | -0.00346 | |
| | (0.0117) | (0.0115) | | | (0.0452) | |
| Male | 0.0271*** | 0.0269*** | | | -0.207*** | |
| | (0.00687) | (0.00684) | | | (0.0315) | |
| NYU | 0.00178 | 0.00110 | | | -0.0812** | |
| | (0.00869) | (0.00865) | | | (0.0386) | |
| Subject fixed effects | No | No | Yes | Yes | No | Yes |
| Obs.ervations | 9,488 | 9,488 | 9,488 | 9,488 | 3,337 | 3,337 |
| R-squared | 0.529 | 0.533 | 0.619 | 0.623 | 0.108 | 0.692 |

Note: Linear probability model for senders. OLS for receivers. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

**Table 4: Summary of absolute mistakes relative to a player's best response**

| | Panel A: Take every round as one period | | | | | |
|---|---|---|---|---|---|---|
| Absolute sender mistake in response to opponent action of: | No Feedback (HBS or NYU) | Feedback (HBS) | p-value (=) | No Feedback (HBS only) | No Feedback (NYU only) | p-value (=, HBS only) |
| last period | 0.2104 | 0.2400 | 0.0043*** | 0.2396 | 0.1940 | 0.9765 |
| Absolute receiver mistake in response to opponent action of: | | | | | | |
| last period | 0.8874 | 0.8267 | 0.0978* | 0.9642 | 0.8402 | 0.0024*** |
| | **Panel B: Take every 5 rounds as one period** | | | | | |
| Absolute sender mistake in response to opponent action of: | No Feedback (HBS or NYU) | Feedback (HBS) | p-value (=) | No Feedback (HBS only) | No Feedback (NYU only) | p-value (=, HBS only) |
| last period | 0.1954 | 0.1802 | 0.1257 | 0.2148 | 0.1843 | 0.0033** |
| Absolute receiver mistake in response to opponent action of: | | | | | | |
| last period | 0.7636 | 0.6860 | 0.0139*** | 0.8426 | 0.7152 | 0.0001*** |

Note: Sender mistake is defined as actual action minus optimal action, where sender action is equal to 1 if report and 0 if not report. Receiver mistake is defined as actual action minus optimal action, where receiver action is equal to the receiver's guess conditional the sender does not report. *** p<0.01; ** p<0.05; * p<0.1.

**Table 5: Regression results on absolute mistakes**

Sender action = report or not
Sender mistake = actual action – best action in response to receiver action in the last period
Receiver action = guess of non-reported number
Receiver mistake = actual guess – best action in response to sender action in the last period
(conditional on the sender does not report)

| | Panel A: Dependent Variable = Absolute Sender Mistake | | | |
|---|---|---|---|---|
| Period definition | Every round | | Every 5 rounds | |
| Feedback | 0.0173 | (dropped) | 0.0174 | (dropped) |
| | (0.0225) | | (0.0252) | |
| Period | -0.000949*** | -0.00102*** | -0.0295*** | -0.0294*** |
| | (0.000352) | (0.000348) | (0.00187) | (0.00186) |
| Period * Feedback | -0.000290 | -7.36e-05 | -0.00848** | -0.00839** |
| | (0.000800) | (0.000791) | (0.00387) | (0.00388) |
| Subject demographics | Yes | No | Yes | No |
| Subject fixed effects | No | Yes | No | Yes |
| Observations | 9,488 | 9,488 | 9,488 | 9,488 |
| R-squared | 0.106 | 0.193 | 0.125 | 0.207 |

| | Panel B: Dependent Variable = Absolute Receiver Mistake | | | |
|---|---|---|---|---|
| Period definition | Every round | | Every 5 rounds | |
| Feedback | 0.00897 | (dropped) | 0.0403 | (dropped) |
| | (0.0800) | | (0.0804) | |
| Period | -0.00273** | -0.00263** | -0.00424 | -0.00574 |
| | (0.00123) | (0.00108) | (0.00605) | (0.00427) |
| Period * Feedback | -0.00576* | -0.00973*** | -0.0323** | -0.0358*** |
| | (0.00298) | (0.00268) | (0.0130) | (0.0113) |
| Subject demographics | Yes | No | Yes | No |
| Subject fixed effects | No | Yes | No | Yes |
| Observations | 3,023 | 3,023 | 2,890 | 2,890 |
| R-squared | 0.030 | 0.458 | 0.034 | 0.649 |

Note: Robust standard errors. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

**Table 6: Summary of player actions (5 states and 10 states)**

| | | | Panel A: Sender Report? (0/1) | | |
|---|---|---|---|---|---|
| Sender's draw (5 states) | Sender's draw (10 states) | Frequency (5st+10st) | No Feedback (5 states) | No Feedback (10 states) | p-value (5st=10st) |
| 1 | 1 | 1495+216 | 7.89% | 12.96% | 0.0004*** |
| | 2 | 189 | | 22.22% | |
| 2 | 3 | 1536+189 | 40.69% | 41.21% | 0.8438 |
| | 4 | 185 | | 71.35% | |
| 3 | 5 | 1479+199 | 83.23% | 83.58% | 0.8599 |
| | 6 | 187 | | 89.84% | |
| 4 | 7 | 1499+185 | 95.86% | 95.85% | 0.9933 |
| | 8 | 173 | | 95.95% | |
| 5 | 9 | 1454+201 | 95.60% | 96.47% | 0.4538 |
| | 10 | 177 | | 98.31% | |
| Obs. | | 9353 | 7463 | 1890 | |

| | | | Panel B: Receiver Guess | | |
|---|---|---|---|---|---|
| Sender's message | Frequency (5 states) | Frequency (10 states) | No Feedback (5 states) | No Feedback (10 states) | p-value (5st=10st) |
| 1 | 118 | 28 | 1.38 | 1.43 | 0.8153 |
| 2 | 625 | 42 | 2.15 | 2.19 | 0.6456 |
| 3 | 1231 | 82 | 3.05 | 3.24 | <0.0001*** |
| 4 | 1437 | 132 | 4.01 | 4.20 | <0.0001*** |
| 5 | 1390 | 168 | 4.91 | 5.24 | <0.0001*** |
| 6 | | 168 | | 6.20 | |
| 7 | | 185 | | 7.12 | |
| 8 | | 166 | | 8.01 | |
| 9 | | 164 | | 8.96 | |
| 10 | | 174 | | 9.91 | |
| Blank | 2662 | 581 | 2.16 | 3.42 | |
| Actual | | | 1.69 | 2.62 | |
| Obs. if reported | 4801 | 1309 | | | |
| Obs. if not reported | 2662 | 581 | | | |
| Obs. total | 7463 | 1890 | | | |

Note: *** p<0.01; ** p<0.05; * p<0.1.

**Table 7: Summary of stated beliefs about sender strategies and Obs.erved reporting rates**

| | No Feedback | Feedback | p-value (=) |
|---|---|---|---|
| Draw | Stated belief of % senders that will report | | |
| 1 | 8.633% | 12.222% | 0.1561 |
| 2 | 27.675% | 31.989% | 0.2252 |
| 3 | 57.333% | 66.067% | 0.0236** |
| 4 | 79.775% | 83.578% | 0.3385 |
| 5 | 86.867% | 89.567% | 0.4762 |
| Implied guess from reported belief | 1.952 | 1.888 | 0.4072 |
| Sender report | % of a specific draw observed in all rounds as receiver | | |
| 1 | 2.269% | 2.301% | 0.9469 |
| 2 | 8.795% | 10.138% | 0.1381 |
| 3 | 14.720% | 16.727% | 0.0754* |
| 4 | 18.489% | 18.047% | 0.6871 |
| 5 | 18.469% | 19.565% | 0.3159 |
| Blank | 37.258% | 33.222% | 0.0095*** |
| Should guess based on observed history | 1.701 | 1.353 | 0.0677* |
| Average actual guess | 2.289 | 1.862 | 0.0002*** |
| Average actual guess for rounds<=25 | 2.420 | 2.064 | 0.0026*** |
| Average actual guess for rounds>25 | 2.152 | 1.541 | 0.0000*** |
| corr(implied guess, should guess) | 0.103 | 0.207* | |
| corr(implied guess, actual guess in the 2nd half) | 0.5859*** | 0.3677*** | |
| corr(should guess, actual guess in the 2nd half) | 0.074 | 0.114 | |

Note: ***p<0.01; **p<0.05; *p<0.1.

**Table 8: Receiver learning from experience (all conditional on senders not reporting the draw)**

| | Dependent Variable = Guess of blank report – last guess of blank report | | | | | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Constant | 0.0615** | 0.0272 | 0.0886** | 0.172*** | 0.145*** | 0.309*** |
| | (0.0288) | (0.0357) | (0.0361) | (0.0408) | (0.0497) | (0.0482) |
| Feedback | 0.134* | 0.144 | 0.0205 | (dropped) | (dropped) | (dropped) |
| | (0.0795) | (0.0915) | (0.116) | | | |
| Round | 0.00162** | 0.00172** | 0.00101 | 0.00190 | 0.00198 | 0.000460 |
| | (0.000820) | (0.000823) | (0.000820) | (0.00133) | (0.00133) | (0.00128) |
| Round * Feedback | -0.00310 | -0.00309 | -0.00325 | -0.00377 | -0.00381 | -0.00629* |
| | (0.00230) | (0.00230) | (0.00256) | (0.00346) | (0.00353) | (0.00334) |
| Over-guessed last time | -0.282*** | -0.250*** | | -0.431*** | -0.401*** | |
| | (0.0313) | (0.0351) | | (0.0395) | (0.0491) | |
| Over-guessed last time * feedback | -0.304*** | -0.314*** | | -0.354*** | -0.376*** | |
| | (0.0972) | (0.108) | | (0.0951) | (0.116) | |
| Under-guessed last time | | 0.0568 | | | 0.0526 | |
| | | (0.0391) | | | (0.0453) | |
| Under-guessed last time * feedback | | -0.0177 | | | -0.0391 | |
| | | (0.0680) | | | (0.0996) | |
| Magnitude of over-guess last time | | | -0.245*** | | | -0.476*** |
| | | | (0.0289) | | | (0.0375) |
| Magnitude of over-guess last time * feedback | | | -0.0963 | | | -0.247*** |
| | | | (0.114) | | | (0.0846) |
| Magnitude of under-guess last time | | | 0.0326 | | | 0.0268 |
| | | | (0.0252) | | | (0.0283) |
| Magnitude of under-guess last time * feedback | | | 0.0728* | | | 0.0473 |
| | | | (0.0438) | | | (0.0479) |
| Subject FE | No | No | No | Yes | Yes | Yes |
| Observations | 2,915 | 2,915 | 2,915 | 2,915 | 2,915 | 2,915 |
| R-squared | 0.064 | 0.065 | 0.125 | 0.128 | 0.129 | 0.276 |

Note: Robust standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

**Table 9: Summary of player actions before and after information on aggregate reporting**

**Panel A: Sender reports or not (0/1)**

| Sender's draw | Frequency | Feedback (round 41-45) (A) | Disclosure rate after feedback (B) | p-value (A = B) | No feedback (rounds 41-45) (C) | Average report after no feedback (D) | Disclosure rate after no feedback (E) | p-value (C = D) | p-value (C = E) |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 386 | 4.35% | 2.27% | 0.589 | 6.04% | 0.00% | 6.45% | 0.249 | 0.895 |
| 2 | 341 | 66.67% | 79.17% | 0.202 | 56.44% | 50.00% | 55.00% | 0.644 | 0.832 |
| 3 | 331 | 100.00% | 95.56% | 0.143 | 89.02% | 92.86% | 81.67% | 0.658 | 0.148 |
| 4 | 332 | 97.87% | 95.65% | 0.550 | 97.52% | 100.00% | 97.06% | 0.617 | 0.844 |
| 5 | 353 | 97.92% | 100.00% | 0.353 | 95.42% | 100.00% | 95.75% | 0.385 | 0.906 |
| Obs. | 1743 | 225 | 225 | | 823 | 75 | 395 | | |
| Average draw if no report | 555 | 1.328 | 1.351 | 0.864 | 1.542 | 1.310 | 1.571 | 0.167 | 0.756 |

**Panel B: Receiver guess conditional on sender not reporting**

| Sender's message | Frequency | Feedback (round 41-45) (A) | Disclosure rate after feedback (B) | p-value (A = B) | No feedback (rounds 41-45) (C) | Average report after no feedback (D) | Disclosure rate after no feedback (E) | p-value (C = D) | p-value (C = E) |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 20 | 1 | 3 | NA | 1.455 | NA | 1.25 | NA | 0.6305 |
| 2 | 205 | 2.083 | 2.039 | 0.5632 | 2.12 | 2.071 | 2.023 | 0.7623 | 0.1364 |
| 3 | 299 | 3.052 | 3.012 | 0.5185 | 3.051 | 3 | 3.041 | 0.5833 | 0.8334 |
| 4 | 323 | 4.011 | 4.034 | 0.5104 | 4.029 | 4 | 4.053 | 0.7011 | 0.5127 |
| 5 | 341 | 4.915 | 5 | 0.3474 | 4.908 | 5 | 4.933 | 0.4406 | 0.6778 |
| Blank | 555 | 1.431 | 1.663 | 0.2423 | 2.011 | 1.759 | 2.114 | 0.1392 | 0.322 |
| Obs. if reported | 1188 | 167 | 168 | | 552 | 46 | 255 | | |
| Obs. if not reported | 555 | 58 | 57 | | 271 | 29 | 140 | | |
| Obs. total | 1743 | 225 | 225 | | 823 | 75 | 395 | | |