

Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure

Ginger Zhe Jin
Michael Luca
Daniel Martin

Working Paper 15-078



Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure

Ginger Zhe Jin
University of Maryland

Michael Luca
Harvard Business School

Daniel Martin
Northwestern Kellogg School of Management

Working Paper 15-078

Copyright © 2015, 2016, 2017 by Ginger Zhe Jin, Michael Luca, and Daniel Martin

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure

Ginger Zhe Jin
University of Maryland and NBER

Michael Luca
Harvard Business School

Daniel Martin
Northwestern Kellogg School of Management

First version: April 2015
Latest version: November 2017

Abstract

This paper uses laboratory experiments to directly test a central prediction of disclosure theory: that strategic forces can lead those who possess private information to voluntarily provide it. In a simple two-person disclosure game, we find that senders disclose favorable information, but withhold less favorable information. The degree to which senders withhold information is strongly related to their stated beliefs about receiver actions, and their stated beliefs are accurate on average. Receiver actions are also strongly related to their stated beliefs, but receiver actions and beliefs suggest they are insufficiently skeptical about non-disclosed information. As a result, senders increase expected returns by strategically withholding unfavorable information, in contrast with classic theoretical predictions.

JEL: D82, D83, C92, L15

Keywords: Communication games, Disclosure, Unraveling, Experiments

Acknowledgements: Patrick Rooney provided excellent research assistance. All errors are ours. Part of the revision was carried out while Jin took leave at the Federal Trade Commission. The opinions expressed here are those of the authors and not necessarily those of the Federal Trade Commission or any of its Commissioners.

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?

A central result in information economics is that market forces can lead 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 result is simple: consumers will treat all non-disclosing companies the same, so the best businesses among those that do not disclose 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 from non-disclosure. In other words, no news is bad news.

The policy relevance of the unraveling result is clear. Information can be important for markets to function properly, and this result suggests that voluntary disclosure could solve asymmetric information problems across a variety of domains. Moreover, this result points to policies that should increase the extent of disclosure. For example, some cities send hygiene scorecards that restaurants can voluntarily post on their doors, which is an attempt to make disclosure both low cost and verifiable. Online labor markets such as Upwork allow users to choose to take exams to verifiably reveal their skillsets.

Yet, the unraveling logic rests on strong assumptions around the ability of consumers to make inferences about a business's decision to withhold information. When customers do not understand the implications of shrouded or missing information, recent theoretical results predict that unraveling can fail to occur (Eyster and Rabin 2005, Gabaix and Laibson 2006, Mullainathan, Schwartzstein, and Shleifer 2008, Heidhues, Koszegi, and Murooka 2016).

In practice, voluntary disclosure is far from complete (Mathios 2000, Luca and Smith 2015, Bederson et al forthcoming). Most restaurants do not post their hygiene grades unless required to, even when they are sent scorecards. Universities often publicize rankings only in which they do well. And most grocery store food items did not include nutritional information until it was mandated. However, because multiple factors can lead to these failures of voluntary disclosure, it is difficult to cleanly test the unraveling prediction and the role of consumer inferences about non-disclosure.

The goal of our paper 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 these tensions. 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 In many markets, such as those with truth-in-advertising laws, firms choose whether or not to reveal information, but can only disclose verifiable information. By prohibiting dishonest reporting, we mirror this feature, and also reproduce the assumptions underlying the unraveling prediction.

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 (a whole number between 1 and 5), 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 from which states are drawn, 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, 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 their information.

Examining the choices of 324 experimental subjects in our main sessions, we find high rates of disclosure among senders with favorable draws. Across treatments, roughly 95% of senders with a random draw of 4 or 5 choose to disclose it. However, senders with lower draws often hide their information from receivers. While the unraveling result predicts that all senders with a draw of 2 will reveal it, we find that the majority of senders with a draw of 2 choose to withhold it.² In fact, withholding information about a draw of 2 maximizes the expected returns of senders because receivers tend to overestimate the draw when senders withhold it.

To help understand what drives this behavior, we elicit beliefs about average receiver guesses and aggregate sender strategies. From this, we learn that sender disclosure decisions are strongly related to their stated beliefs about receiver guesses of non-

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).

² As a consequence, the decision to withhold is not perfectly revealing of the state. The variance in the state when it has not been disclosed is almost a full unit (0.947).

disclosed states, and that receiver guesses are strongly related to their stated beliefs about sender strategies. However, while sender beliefs are correct on average, receiver beliefs are not. Receivers correctly estimate that almost no senders disclose draws of 1, but they underestimate the rates of disclosure among senders with higher draws, particularly those with draws of 2 or 3. This implies that they believe a non-disclosed draw is more likely to be 2 or 3 than in reality. In other words, receivers are insufficiently skeptical about undisclosed information – the extent to which no news is bad news.

Our findings are consistent with 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 Lovo (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.³ Luca and Smith (2015) show that MBA programs with lower rankings are less likely to display them on their website.

Our results complement these findings from the field. Existing field studies provide important evidence on the extent of, and barriers to, disclosure in different contexts. However, it is difficult to directly test consumer inferences and the unraveling hypothesis in the field because a myriad other factors can influence disclosure and consumer decisions in each setting. In contrast, the lab allows us to focus directly on the behavioral assumptions and beliefs underlying unraveling, but it is difficult to generalize the magnitude of the effects we observe to different field settings, precisely because there are many other factors at play in the field.

One important source of field variation that we can study in the lab is the impact of experience and opportunities for learning on unraveling. We have participants play our disclosure game 45 times, and we exogenously vary the extent of the feedback we provide and whether subjects switch roles during the experiment. In the benchmark treatment, subjects receive no feedback, so receivers only learn the true state when the sender reveals it, and subject roles are fixed for the entire experiment. In a second treatment, subjects also receive no feedback, but are randomly assigned roles in each round, so they experience both roles. While opportunities for learning, either through disclosure or introspection, do exist in these two treatments, we substantially increase the opportunities for learning in a third treatment. In that treatment, participants are debriefed after each round and are also randomly assigned a role in each round. In the debriefing

³ Brown, Camerer, and Lovo (2013) demonstrate how data on movie openings can be used to differentiate between equilibrium and non-equilibrium behavior (specifically related to the extent that naïveté limits unraveling in settings of verifiable disclosure).

after each round, they are told the true state, whether it was revealed by the sender, and the receiver's guess. In addition, participants from the second and third treatments are given aggregate information toward the end of the session about the overall likelihood of disclosure for each draw.

For the first treatment (“no feedback & fixed role”) and the second (“no feedback & random role”), behavior does not converge to the unraveling equilibrium during the 45 rounds of our experiment. However, in the third treatment (“feedback & random role”), behavior ultimately converges to the unraveling equilibrium. Moreover, using regression analysis we find little evidence of decreasing guesses or increasing disclosure of draws of 2 for the first treatment. We also find little evidence of a difference between the first and second treatments along these dimensions, which suggests that role switching (alone) is not enough for players to fully recognize the strategic element behind non-disclosure. However, we find strong evidence of a difference along these dimensions between the first treatment and the third. This suggests that feedback is an important factor in producing convergence in our experiment. Yet, when aggregate feedback about disclosure rates is given in later rounds of the second treatment, it does not lead to behavior fully consistent with the unraveling prediction, which suggests that it is important that the feedback is immediate, direct, and/or repeated.

Overall, our results help to shed light on the economics of voluntary disclosure. In situations where immediate and direct feedback about non-disclosed information is limited, our findings suggest there is reason to be skeptical that full unraveling will occur. One reason for limited feedback could be infrequent transactions, as in the markets for house or car purchases. Even in markets with frequent transactions, feedback can be limited when consumers are inattentive to or unable to process feedback about non-disclosed information. For example, when restaurants choose not to disclose their hygiene ratings or calorie counts, it can be difficult or time-consuming for consumers to assess this information about the restaurant immediately after having completed their meal.

This paper provides two main contributions, which are discussed in more detail in the next section. First, we complement existing empirical studies from the lab and field by providing evidence of strategic naïveté in verifiable disclosure using a novel and clean laboratory experiment.⁴ We also elicit beliefs from both senders and receivers, which provides more direct evidence on the underlying mechanisms.⁵

⁴ Our relationship to existing verifiable disclosure experiments is proved 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

Second, we vary the feedback provided to subjects, and explore its role in producing convergence to equilibrium behavior. 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, direct, and repeated feedback, consumers may not engage in behaviors that induce full disclosure. The type of aggregate feedback that is more common in the field and in information interventions does not appear to change behavior in our experiment.

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, and Section 5 reports the results of our experiments. A brief discussion of the results is offered in Section 6.

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 play 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

hypothetical thinking, which has been shown to be difficult in voting games (Esponda and Vespa 2014).

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

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, 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 about 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 receiver naïveté in voluntary disclosure through a controlled laboratory experiment where beliefs are elicited. Second, such naïveté is not easily eliminated, even if receivers are provided information about aggregate disclosure behavior and have played as senders for some rounds. 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 Sapro 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 Sapro (2003) complement these findings by also showing what happens when there is a 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) study the use of both vague and ambiguous messages, and find that an announcer in coordination games might want to use such messages.

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 contrast fixed roles with role switching, to vary the feedback provided to subjects, to elicit beliefs about receiver guesses and sender strategies, or to provide information about aggregate sender behavior.

2.3 Beliefs and Play in Games

Central to any strategic interaction is the set of beliefs that people hold about each other. This has given rise to work in the experimental economics literature based on the following question: Do people hold correct beliefs about how other people play and do they best respond to these beliefs?

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 receiver guesses and sender strategies. This allows us to compare stated beliefs across treatment arms and also to see whether subjects are acting in line with their beliefs. We find a strong, positive, and statistically significant correlation between beliefs about receiver guesses and the probability of the sender reporting the underlying state. Likewise, 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. Furthermore, we demonstrate asymmetry between the correctness of sender beliefs and the correctness of receiver beliefs. While sender beliefs are correct on average, average receivers are insufficiently skeptical about undisclosed information.

This asymmetry motivates senders to withhold non-favorable information more than theory predicts.

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 b (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 B , 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 b is just $M(b)=\{b,null\}$.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 B . 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 $U_R(a, b)$, which is concave in the receiver's action a and reaches its maximum when a is equal to b . In other words, the receiver benefits more from selecting an action that is closer 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 B), and if there is no report, the receiver takes the action that is the minimum element in B .⁸ In other words, the sender always reports his or her type (unless it is the worst

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 confounding factors.

⁸ In the appendix, we show that sender strategies in all other Bayesian Nash equilibria of this game cannot satisfy a condition that is clearly consistent with the behavior we observe. As a consequence, subject choices are not in line with other Bayesian Nash equilibria of this game.

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 B , the sender is indifferent between reporting or not, so any mixture over these actions is consistent with equilibrium.

4 Experimental Design

In our experiments, subjects completed 45 rounds and then, depending on the session, specific additional tasks. Subjects were told at the beginning of the experiment that they would complete additional tasks, but were given no details about the tasks they would face later. The appendix contains the full set of instructions given before the start of the experiment. Instructions for an additional task were presented to subjects on the computer screen just before the start of a task.⁹

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 converted into U.S. dollars at a treatment and role-specific rate, and their final payment was rounded up to the nearest non-negative whole dollar amount.

4.1 Main Sessions

Our main sessions were conducted at the Computer Lab for Experimental Research (CLER) facility at the Harvard Business School (HBS). In this laboratory, subjects are separated with dividers, and each subject was provided with a personal computer terminal.

4.1.1 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 7.7%.

⁹ Our experiment was programmed and run using the z-Tree software package (Fischbacher 2007).

In each round and for each pairing, one subject was the sender, and the other subject was the receiver. 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 $B=\{1,2,3,4,5\}$. 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 that if the area for messages on the receiver’s screen was left blank, the instructions were clear that 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 was $A=\{1,1.5,2,2.5,3,3.5,4,4.5,5\}$.¹⁰ There was also no time limit for receiver decisions.

Receiver payoffs in each round were $ECU_R = 110 - 20|b - a|^{1.4}$, where b is the secret number and a is the receiver’s guess.¹¹ These payoffs decrease monotonically as the guess moves further from 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 increase monotonically with receiver guesses because guesses cannot 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.

¹⁰ The action space of receivers was made sufficiently rich that the unique sequential equilibrium involves full unraveling.

¹¹ We allowed subjects accrue ECU in all rounds because payoffs could vary substantially between roles and realizations of the state, and we wanted performance to play a larger role than luck in final payments. Cai and Wang (2006) use similar payoff functions and also paid subjects every round. However, this approach introduces the possibility of wealth and portfolio effects. To ameliorate such effects, subjects were not told the cumulative payoffs they had earned so far in the experiment.

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 in the instructions about these two features of sender and receiver payoffs.

4.1.2 Treatment Variation

Our primary treatment variations occurred along two dimensions. First, we varied the information provided as feedback to subjects after each round (no feedback vs. feedback), and second, we varied the way that roles were assigned (fixed role vs. random role). Both sources of variation were used to study the channels through which subjects learn in this setting. We ran three combinations: “no feedback & fixed role”, “no feedback & random role”, and “feedback & random role”.

In our “no feedback” treatments, 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.

In our “fixed role” treatment, subjects were randomly assigned a role at the beginning of the session, and they stayed in that role throughout the entire experiment. Instead, in our “random role” treatments, subjects were randomly assigned roles before each round, so that roles might change after each round. In both cases, a subject was equally likely to be assigned either role. As a result, the likelihood of a subject experiencing both roles by round 5 in the “random role” treatments was 93.75%. In the random role treatment, ECU were converted into U.S. dollars at a rate of 200 to 1, but to equalize expected payments across subjects in the fixed role treatment, ECU were converted at a rate of 150 to 1 for senders and 250 to 1 for receivers.

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. In addition, between rounds subjects only received feedback about their pairing, not all pairings in the session.

4.1.3 Belief Elicitation and Additional Tasks

After completing 45 rounds, subjects were asked to guess both the rate at which senders reported each secret number and the average receiver guess of non-reported secret numbers during the preceding 45 rounds. The purpose of these questions was to assess whether subject beliefs about sender strategies influenced their decisions as receivers and whether subject beliefs about receiver guesses influenced their decisions as senders. These guesses were not incentivized, but in a recent paper, Trautmann and Kuilen (2015) show that such “introspective” elicitation can yield accurate beliefs.

After beliefs were elicited, subjects completed an additional task. In our main sessions, subjects either completed a “high incentives” task or an “aggregate feedback” task.

The “high incentives” task was designed to better understand the impact of incentives on mistakes. This task mirrored an earlier choice, but with much higher incentives. In this task, subjects played once more in the role of sender or once more in the role of receiver, but in both cases they played against a computer instead of a human, and the computer played a strategy that was designed to mimic the unobserved decision of a previous opponent. Subjects were not reminded of the choices they had made previously.

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 decisions, 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.

Importantly, for this decision the payoffs of the subject were ten times the rate in the initial 45 rounds. This design allows us to hold fixed the strategy of the opponent, so that the impact of incentives can be isolated to just one side of the pairing. Niederle and Vesterlund (2007) use a similar approach to hold fixed opponent strategies.¹²

¹² There are two potential confounds for this task. First, in this choice, there are no payoff consequences for their opponent, so social preferences related the opponent’s payoffs are no longer in play. Second, the random round could be drawn from any part of the experiment, so if there is a large time trend in behavior, the subject may choose differently because of additional uncertainty over actions.

In addition, to better understanding the interaction between beliefs and information, we used a second additional task in which subjects were shown information about the play of all subjects in the first 45 rounds, guessed again a non-disclosed secret number from a previous round, 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.

We call this task the “aggregate feedback” task because subjects were shown the number of times that each secret number was reported and not reported for all subjects during the first 45 rounds of their session. This provided enough information to determine both the average reported secret number and the average non-reported secret number.

4.2 10-State Robustness Sessions

To get a sense for how the size of the state space might impact our findings, we ran an additional robustness sessions with 10 secret numbers, where $B = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$, which is twice as large as the state space in the main sessions. Here again we allow receivers to guess half-unit intervals, so the action space is $A = \{1, 1.5, 2, 2.5, \dots, 9, 9.5, 10\}$.

To keep payoffs in a comparable range to the main sessions, the distance from the ideal action is divided in half in the payoff functions, so that receiver payoffs are $ECU_R = 110 - 20|(b - a)/2|^{1.4}$ and sender payoffs are $ECU_S = 110 - 20|(10 - a)/2|^{1.4}$. As a result, in these robustness sessions the payoffs for senders and receivers when the receiver guesses a and the secret number is b is the same as when the receiver guesses $a/2$ and the secret number is $b/2$ in the main sessions.

Aside from increasing the set of secret numbers and changing the payoff table, the experimental design and instructions are the same as in the main sessions. In order to ensure sufficient statistical power, we did not vary treatments in the robustness sessions – all subjects completed the “no feedback & random role” treatment. We also conducted these robustness sessions in the CLER facility at HBS.

5 Results

This section examines sender and receiver behavior in both our main sessions and robustness sessions. To provide a complete picture of behavior, we look both at choices pooled across rounds and how choices evolve from round-to-round.

5.1 Subjects in the Main Sessions

In our main sessions, we observed 324 subjects complete a total of 14,580 decisions, which corresponds to 7,290 pairings. Over 23 sessions, the mean session size was 14.1. We used a show-up fee of \$5, and on average subjects earned \$26.60. The minimum payment was \$14 and the maximum payment was \$37.

We assigned 114 subjects to “no feedback & fixed role” sessions, 120 subjects to “no feedback & random role” sessions, and 90 subjects to “feedback & random role” sessions. All subjects in the no feedback & fixed role sessions completed the “high incentives” additional task, and all subjects in the no feedback & random role and feedback & random role sessions completed the “aggregate feedback” additional task.

Table 1 presents summary statistics for all 324 subjects in our main session. For 27.8% of these subjects, feedback was provided after each round, which corresponds to our feedback treatment, and for 64.8% of subjects, roles were randomly determined in each round, which corresponds to our random role treatments. In terms of demographics, we had roughly an equal number of men and women, and a large majority of subjects were undergraduates and native English speakers. Around 15% of subjects reported having a friend present in the room during the session. In the regressions presented in this paper, we either control for these demographic factors or use subject fixed effects.

5.2 Sender Disclosures and Receiver Guesses

When pooling across rounds, our primary qualitative findings are that senders disclose favorable states more often than less favorable states and that receivers tend to over-guess non-disclosed states, and these features of the data are summarized by Table 2.

Looking first at senders, the average reporting rate is above 80% when the draw is equal to the average state (a secret number of 3) and above 90% when the draw is 4 or 5. For lower draws, however, the average reporting rate drops to 44.5% for draws of 2 and 11.0% for draws of 1. The theoretical predictions are a reporting rate of 100% for draws of 2 and anywhere between 0% and 100% for draws of 1.

Between treatments, the sender reporting rate differs the most for draws of 2. When the draw is 2, the reporting rate is not significantly different (for a two-tailed test of proportions) between the no feedback & fixed role treatment and the no feedback & random role treatment (p -value=0.840), but is significantly different between either of those treatments and the feedback treatment (p -value=0.0117 for the fixed role treatment and p -value=0.0185 for random role treatment). This difference is reflected in the average

secret number when the secret number is not reported. When there is no disclosure, the average secret number is smaller in the feedback treatment than in the no feedback treatments, though it is only statistically significant (for a two-tailed t-test) between the no feedback & random role and feedback & random role treatments (p -value= 0.0148).

Table 2 also presents average receiver guesses by treatment, conditional on whether the sender reports 1, 2, 3, 4, 5, or nothing. 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). One explanation for both effects could be that random errors generate a boundary effect. Given that under-guessing at the high extreme decreases with experience and feedback, we believe that this effect is indeed due to errors. However, we believe that a more plausible explanation for over-guessing at the low extreme, relates to the social preferences of the receivers.¹³ Because senders tend to be silent when the draw is close to the low extreme (1, 2), receivers may want to compensate those who report out of feelings of guilt, even if over-guessing reduces their own payoff. This effect is present in all treatments, even when there are ample opportunities to learn.

When senders choose not to disclose, receivers guess 2.243 on average in the no feedback & fixed role treatment and 2.283 on average in the no feedback & random role treatment, and this is not significantly different (p -value=0.3472 for a two-tailed t-test). In the feedback & random role treatment, the average guess is lower (1.897) and this significantly different from the other treatments at the 1% level. There is a similar pattern for the average amount of over-guessing (how far the guess is above the actual secret number). The amount of over-guessing is similar and not significantly different between the two treatments without feedback (p -value=0.6557), and the treatment with feedback is significantly different from the other treatments at the 1% level.

5.2.1 Robustness Check: 10-State Sessions

There are 84 subjects in our 10-state robustness sessions, and as mentioned previously, all subjects in those sessions were assigned to the no feedback & random role treatment. Table 3 provides the summary of player actions in the no feedback & random role treatment for the main sessions and for the 10-state robustness sessions. As when there are 5 secret numbers, the reporting rate increases monotonically with the secret number in the sessions with 10 secret numbers. The reporting rate for a secret number of 3 in the

¹³ The appendix contains an extended discussion of the various ways that social preferences might impact receiver behavior.

robustness sessions is 41.2%, which is comparable to and not statistically different from the reporting rate for a secret number of 2 in the primary study of 42.6% ($p=0.7379$ for a two-tailed test of proportions). In addition, the reporting rate in the robustness sessions 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 main sessions.

A secret number of 3 in robustness sessions and a secret number of 2 in main sessions are also comparable in the sense that a risk neutral sender would not want to report secret numbers of 3 in the robustness sessions (when pooling choices across rounds). The average guess for a non-reported secret number is 3.419 with a 95% confidence interval of 3.254 to 3.584. As in the main sessions, the average guess is above the average actual non-reported secret number in the robustness sessions. In the robustness sessions, 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 when secret numbers are reported. Once again, guesses are higher than reported secret numbers for low secret numbers and lower than reported secret numbers for high secret numbers, though the effect is smaller for high secret numbers.

In short, our primary qualitative findings for sender and receiver behavior in the main sessions – incomplete disclosure by senders and over-guessing of non-reported states by receivers – are robust to enlargement of the state space.

5.3 Departures from the Highest Expected Payoff

To quantify the impact of incomplete disclosure and receiver over-guessing on payoffs, we measure how far a subject is from taking the payoff-maximizing action in each decision problem, which provides a rough sense for the size and consequences of the “mistakes” they are making. To do this, we construct the average opponent strategy from our data, determine the expected payoffs for each possible action, and then calculate how far the expected payoffs for the taken action are from the highest expected payoff for taking any action.¹⁴ For senders, the possible actions are reporting or not reporting the secret number. For receivers, we just consider the guesses that are available to them, which is important because they are limited to guessing half-units.¹⁵

Table 4 reports the monetary losses that result from the actions taken in our main sessions, our robustness sessions, and in our high incentives task. Across these settings,

¹⁴ Because the minimum possible payoff can be negative, we normalized payoffs by the distance from the minimum possible payoff for the realized state.

¹⁵ We thank an anonymous referee for suggesting this method.

we find that receivers are between 9.3% and 13.3% away on average from the highest expected payoff and that senders are between 3.4% and 7.3% away on average from their highest expected payoff. In all settings and roles, the losses are significantly different from zero at a 1% level (using a two-sided t-test). We also show the losses of receivers relative to the payoffs they would make in the unraveling equilibrium. These losses range from 19.9% to 26.4%.

In addition, Table 4 illustrates two more findings. First, senders have smaller losses than receivers in all settings. Second, when incentives are increased tenfold, receivers continue to have similar percentage losses. Receiver losses in the first 45 rounds of the no feedback & fixed role treatment of the main sessions are not significantly different from receiver losses in the high incentives task ($p=0.5510$). This does not appear to be driven by low power, as sender losses are significantly different in the same comparison ($p=0.0260$).

These losses are reflected in the size of receiver mistakes, as measured by the distance from the guess with highest expected payoff (normalized by the size of the state space). Across treatments, receivers are 14.1% to 17.1% away from the guess with the highest expected payoff. Even when payoffs are increased tenfold, receivers are still 14.4% away, and as the state space is stretched in the robustness treatment, receivers are still 15.7% away.

Our calculations take an ex-ante perspective, so when determining the highest expected payoff for receivers, we assume that all states are equally likely to happen and determine the average sender behavior separately for each state. In addition, we pool together all rounds when determining average sender and receiver behavior, which is equivalent to assuming that a subject is equally likely to face an opponent from any round. While this is exactly the way that opponents are determined in the high incentives task, such an assumption may not be suitable for treatments where we observe learning.

To adjust for dynamic considerations, in Table 5 we calculate sender and receiver losses with the assumption that a subject is equally likely to face any opponent from the previous block of 5 rounds. However, the magnitude of the losses does not change dramatically with the change in reference group. Across treatments in the main sessions, average sender losses range from 5.1% to 6.8%, and the average receiver losses range from 9.4% to 12.9%.

5.4 Stated Beliefs: Senders

As described in Section 4, after all 45 rounds were completed and before any additional

tasks were undertaken, we asked subjects in the main sessions to guess the average receiver guess when the secret number was not reported over all 45 rounds. The responses of those who played the role of sender at least once are given in Panel A of Table 5.

Subjects in the fixed role treatment have the highest average guess (2.273), followed by those in the no feedback & random role treatment (2.143), and the feedback & random role treatment (1.617). The first two are similar in size and are not significantly different ($p=0.4330$ for a two-tailed t-test). However, stated beliefs are significantly different between both of no feedback treatments and the feedback treatment at a 1% level. As shown in Panel B of Table 5, these guesses are close to the actual averages (2.243, 2.283, and 1.897), particularly for the no feedback treatments.

We find that sender disclosure decisions are largely consistent with their beliefs about receiver guesses. Regardless of treatment, over 85% of decisions are consistent with reporting if and only if the secret number is below the belief of the average receiver guess of the non-reported secret number. Over the last 15 rounds, this rate rises to 89.1% in the fixed role treatments, to 91.3% in the no feedback & random role treatment, and to 92.1% in the feedback and random role treatment. It may not be surprising that the rate of consistency increases given that beliefs are collected after all 45 rounds are complete.

5.5 Stated Beliefs: Receivers

After all 45 rounds were completed, we also ask subjects who played the role of receiver at least once to guess the percentage of senders who reported each secret number over all 45 rounds. The frequency of their responses is given in the three panels of Figure 1, along with the average guesses and the actual rates at which each secret number was reported.

Larger bubbles correspond to a larger number of guesses in an interval of 5 percentage points, and the thick 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 upper lines, which represent the actual reporting rate, pass through this region, so these higher guesses represent accurate beliefs. The thin line represents the reporting rate for the last 15 rounds, and the medium line represents the average reporting rate for all 45 rounds.

Regardless of the treatment, average guesses for reporting rates of draws of 2 and 3 are well below the actual averages, and they are significantly different at a 1% level for all treatments (for a two-tailed t-test). Although the differences in this gap between

treatments are not significant, the feedback treatment has the largest gap between average guess and actual average for draws of 2 (18.3%).

Because receiver beliefs of sender disclosure strategies are significantly different from the actual disclosure rate of an average sender, this suggests that receivers are not sufficiently skeptical about non-disclosure. More importantly, this misbelief is present in all treatments and was reported after subjects have played 45 rounds. This implies that repeated play might have a limited role in correcting the mistaken beliefs of receivers.

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, the gap between average guesses and actual rates increases if we just consider actual rates in the last 15 rounds.

Comparing treatments, the average guess of the disclosure rate for draws of 2 is very similar between the no feedback treatments (26.8% for fixed roles and 27.7% for random roles). The average guess is higher for the feedback treatment (32.0%), but this is not significantly different from the no feedback treatments (p-value=0.2113 for fixed roles and p-value= 0.2252 for random roles).

5.6 Round-by-Round Dynamics: Sender Disclosures

Next, we analyze the dynamics of behavior across 45 rounds and with varying levels of feedback, starting with sender behavior.

Figure 1 shows the overall sender reporting rate for each treatment by block of 5 rounds, and Figure 2 shows the sender reporting rate for each treatment for each of the first three secret numbers separately. As in Figure 1, the dotted line represents the feedback treatment and the solid line represents the fixed role treatment. Without controlling for any other factors, it appears that there is an increasing trend for draws of 2 and 3 for the feedback treatment and possibly for the other treatments as well.

Table 6 investigates these trends using regression analysis with demographic controls and session fixed effects or subject fixed effects. When not using subject fixed effects, we also include a dummy variable for whether the draw is above the sender's guess of the average receiver guess of non-disclosed numbers.

In specifications 1 and 2, all draws are included in the analysis, but separate dummy variables are provided for each of the draws. In keeping with our previous findings, the probability of reporting increases monotonically with the draw and the difference

between a draw of 1 and any other draw is statistically significant at the 1% level holding round, draw, and individual characteristics fixed.

In terms of dynamics, we find that the reporting rate is much lower in the first 5 rounds, and this effect is statistically significant at the 5% level for both specifications. There appears to be a small positive time trend on the reporting rate even for the no feedback treatment, but this effect is only significant at the 10% level. There also appears to be a positive and significant difference between treatments in the time trend, but this is only significant at the 10% level between the no feedback & fixed role treatment and the feedback & random role treatment.

If we look just at draws of 2 (specifications 3 and 4), there is not a statistically significant time trend for the no feedback & fixed role treatment, but there is a statistically significant difference between this treatment and the feedback treatment. The effect size is also relatively large for the difference in time trends between these two treatments.

In these regressions, we also see that the impact of beliefs is largely and highly significant. Controlling for the draw itself, along with time trends and individual characteristics, the probability of reporting increases almost 25% if the sender believes that the receiver will guess lower than the draw (taking stated beliefs at face value).

5.7 Round-by-Round Dynamics: Receiver Guesses

Figure 4 plots the average receiver guess for each block of 5 rounds, conditional on senders not reporting. While all three treatments start out at a similar place, the feedback treatment appears to diverge from the others and has a clearer time trend.

Panel B of Table 5 provides more detail on these trends, but again without any controls. The guess of non-reported secret number appears to be decreasing for all treatments, but this is a bit misleading, as the actual secret number also appears to be decreasing for all treatments. If we look instead at the amount of over-guessing, it appears that the feedback treatment has the biggest drop, with almost no over-guessing in the final block of 15 rounds.

Table 7 shows the output of a number of regressions on the dynamics of receiver guesses. All specifications include both a version with individual demographic controls and subject fixed effects. Looking at specifications 1 and 2, there is a large and statistically significant impact guessing in the first 5 rounds: guesses are much higher during those rounds. In the no feedback & fixed role treatment, there is not a statistically significant

time trend otherwise. However, there is a large and statistically significant difference in time trends between this treatment and the feedback treatment.

As with senders, beliefs appear to have a large and highly significant impact on guesses. For receivers, we use their stated beliefs of sender strategies to generate an implied average non-reported secret number using an equal probability of each state and Bayes' Rule. Controlling for time trends and individual factors, we find in specification 1 that there is a very high correlation between these implied beliefs and what subjects guess.

To investigate why receiver guesses decrease over time in the feedback treatment, 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. For specifications 1 and 2, we find strong evidence that subjects who were informed in the feedback treatment that they guessed too high decreased their guesses the next time they had an opportunity to do so. This effect is statistically significant at the 1% level for both specifications.

There are two potential concerns with these results. First, if subjects have access to other sources of learning, we would expect this effect to exist for all treatments, and it does. However, the effect is stronger for subjects receiving feedback. Second, mean reversion from extreme guesses could produce similar results because very high guesses are almost surely over-guesses and very low guesses are almost surely not over-guesses. In specifications 3 and 4, we examine whether the effects we observe in specifications 1 and 2 hold also for intermediate guesses. We find that the effect in the no-feedback treatment diminishes, but the effect in the feedback treatment stays large. It is statistically significant in specification 3 and has a p-value of 0.05 in specification 4.

In addition, the impact of feedback is not symmetric. There is not a statistically significant relationship between receiving feedback and under-guessing for any of the specifications. Moreover, the point estimates of the effect sizes are much smaller than for over-guessing. This asymmetry, and its predominant presence in the feedback treatment, suggests that receivers corrected their over-guesses by learning from round-by-round feedback rather than through mean reversion.

5.8 Learning 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 aggregate feedback additional tasks, which are all subjects in the two random role treatments. 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. These results are provided in Table 9.

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 subjects in the no feedback treatment.

The table also compares the guesses made when senders did not report their secret number. If anything, the information intervention causes guesses to rise on average. However, there is not a significant difference in guesses before and after the informational intervention for either of the treatments.

6 Discussion

Our findings shed light on a fundamental inference problem that prevents full unraveling in voluntary disclosure, and the conditions under which full unraveling is most likely to occur. In contexts with little or no feedback, receivers are not sufficiently skeptical about undisclosed information. In our experiments, these mistakes can persist for the full 45 rounds of the experiment, and as a result, information senders can profit by limiting disclosure. However, round-by-round feedback about mistakes can result in behavior that converges to the predictions based on full unraveling.

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 unless buyers receive fast and precise feedback about mistakes after each transaction, market forces can be insufficient to close the information gap between sellers and buyers.

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 predictions after a 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. However, there is growing evidence that mandating disclosure may not be sufficient for achieving the desired outcomes (Loewenstein, Sunstein, and Golman 2014, Jin, Luca, Martin 2017), which calls for more future research.

7 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.
- Bederson, B., Jin, G., Leslie, P., Quinn, A., & Zou, B. (forthcoming) Incomplete Disclosure: Evidence of Signaling and Countersignaling. *American Economic Journal: Microeconomics*.
- 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. *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.
- Costa-Gomes, M. A., & Weizsacker, 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.
- Dickhaut, J., Ledyard, M., Mukherji, A., & Sapro, 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., & Vespa, E. (2014). Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 6(4), 180–202.
- Eyster, E., & Rabin, M. (2005). Cursed equilibrium. *Econometrica*, 73(5), 1623–1672.
- 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). Shrouded 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. (2016). Inferior products and profitable deception. *Review of Economic Studies*, 84(1), 323–356.
- Hirshleifer, D., & Teoh, S. H. (2003). Limited attention, information disclosure, and financial reporting. *Journal of Accounting and Economics*, 36(1), 337–386.
- Jin, G. Z. (2005). Competition and disclosure incentives: An empirical study of HMOs. *RAND Journal of Economics*, 93–112.
- Jin, G. Z., Luca, M., & Martin, D. (2017). Complex disclosure. Mimeo.
- 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.
- Loewenstein, G., Sunstein, C. R., & Golman, R. (2014). Disclosure: Psychology Changes Everything. *Annual Review of Economics*, 6(1), 391–419.
- 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*

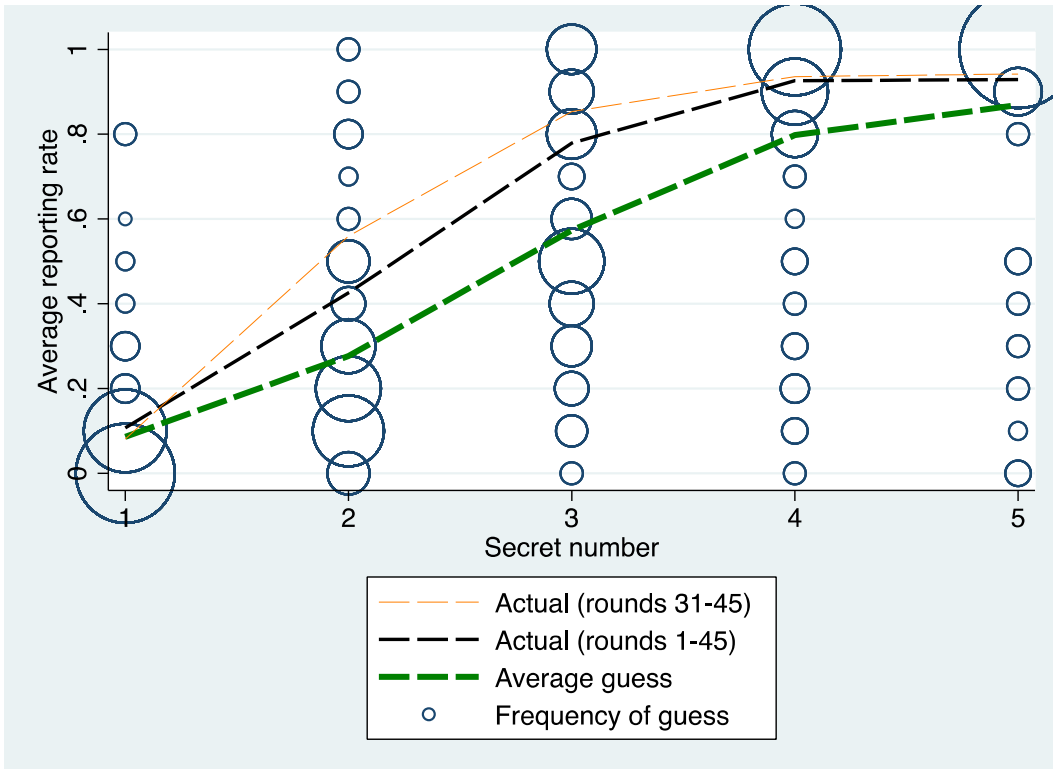
- Economics*, 16(3), 328–340.
- Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *Bell Journal of Economics*, 12(2), 380–391.
- Milgrom, P., & Roberts, J. (1986). Relying on the information of interested parties. *RAND Journal of Economics*, 17(1), 18–32.
- Mullainathan, S., Schwartzstein, J., & Shleifer, A. (2008). Coarse thinking and persuasion. *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.

Figure 1. Sender disclosure rates and guesses of sender disclosure rates (main sessions).

Panel A. No feedback & fixed role treatment.



Panel B. No feedback & random role treatment.



Panel C. Feedback & random role treatment.

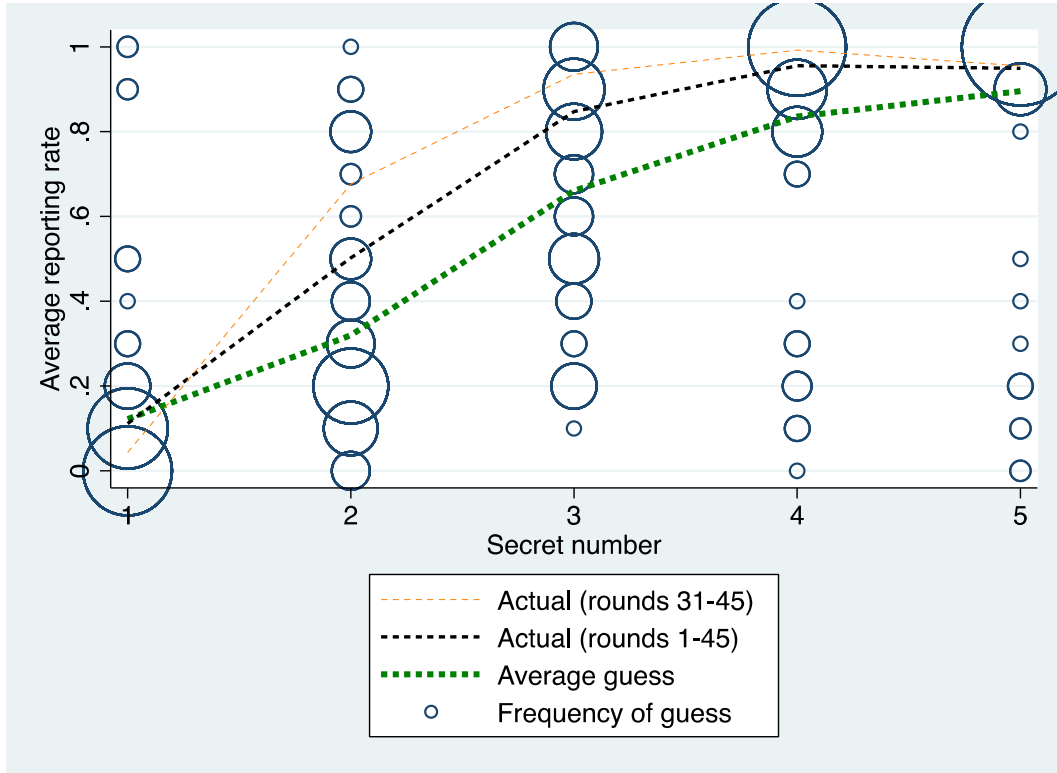


Figure 2. Sender disclosure rates by round (main sessions).

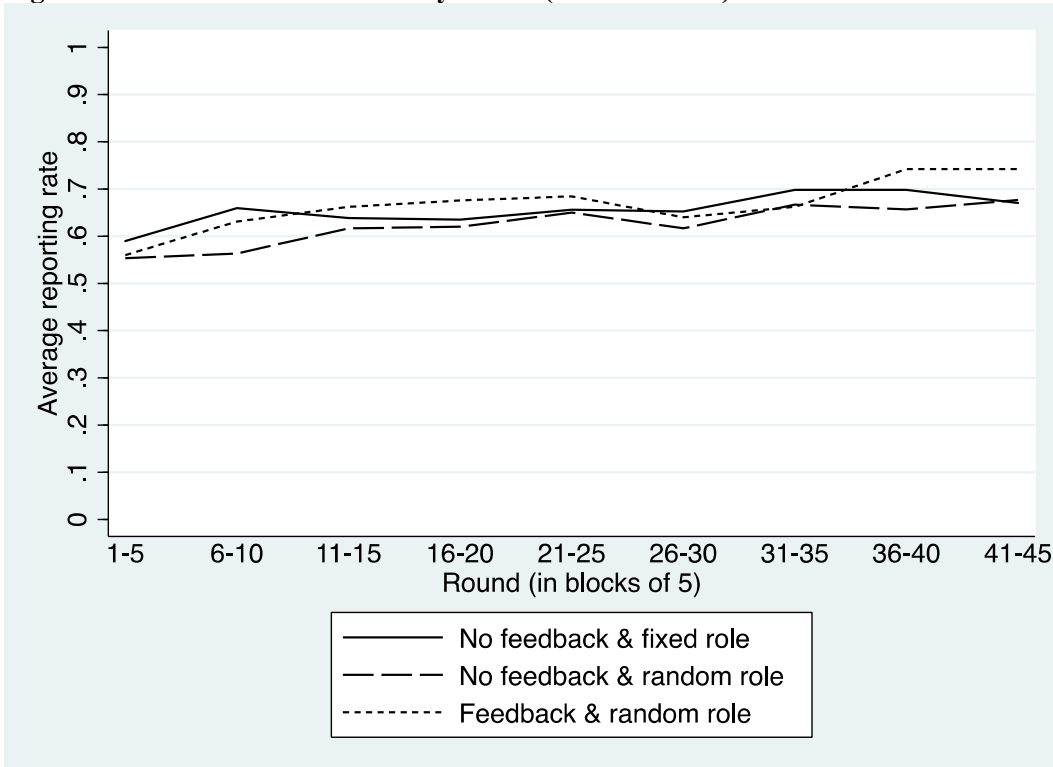


Figure 3. Sender disclosure rates by round for secret numbers of 3, 2, or 1 (main sessions).

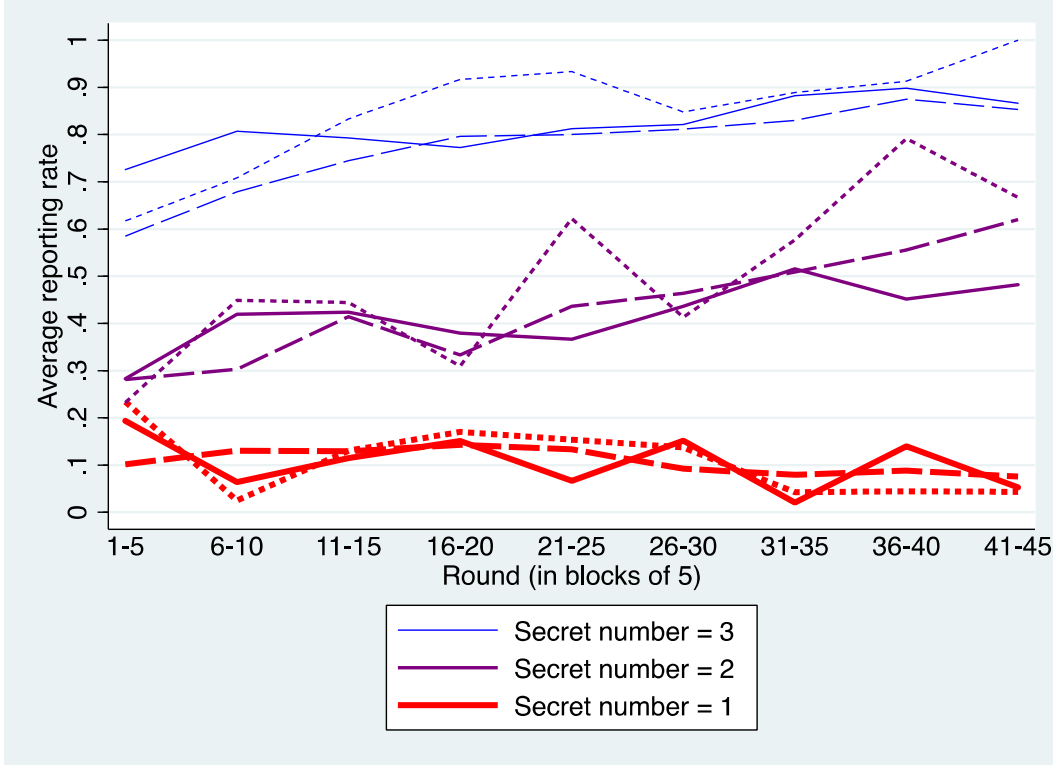


Figure 4. Receiver guesses of non-disclosed secret number by round (main sessions).




Table 1. Summary statistics for main sessions. Observation is per subject. Value is missing if demographic information not provided by the subject.

VARIABLES	N	mean	sd
Number of subjects in the session	324	17.16	7.789
Feedback provided (dummy)	324	0.278	0.449
Random role (dummy)	324	0.648	0.478
Undergraduate (dummy)	324	0.713	0.453
Male (dummy)	324	0.494	0.501
Native English speaker (dummy)	321	0.850	0.357
Friend in the session (dummy)	324	0.145	0.353

Table 2. Summary of player actions in main sessions.

VARIABLES	No feedback Fixed role		No feedback Random role		Feedback Random role	
	N	mean	N	mean	N	mean
Report (secret number=1)	490	0.110	568	0.107	421	0.112
Report (secret number=2)	529	0.420	552	0.426	406	0.502
Report (secret number=3)	533	0.818	507	0.779	400	0.848
Report (secret number=4)	508	0.945	540	0.926	383	0.956
Report (secret number=5)	505	0.968	533	0.929	415	0.949
Secret number (no report)	884	1.734	1,014	1.802	675	1.680
Guess (report=1)	54	1.250	61	1.533	47	1.298
Guess (report=2)	222	2.203	235	2.243	204	2.157
Guess (report=3)	436	3.002	395	3.061	339	3.071
Guess (report=4)	480	3.897	500	4.009	366	4.016
Guess (report=5)	489	4.707	495	4.825	394	4.968
Guess (report=blank)	884	2.243	1,014	2.283	675	1.897
Guess - secret number (no report)	884	0.508	1,014	0.481	675	0.217

Table 3. Summary of player actions in no feedback & random role treatment for main sessions (5 secret numbers) and robustness sessions (10 secret numbers).

VARIABLES	Main sessions		Robustness sessions	
	No feedback		No feedback	
	Random role		Random role	
	N	mean	N	mean
Report (1-baseline or 1-robustness)	568	0.107	216	0.130
Report (2-robustness)			189	0.222
Report (2-baseline or 3-robustness)	552	0.426	199	0.412
Report (4-robustness)			185	0.714
Report (3-baseline or 5-robustness)	507	0.779	201	0.836
Report (6-robustness)			187	0.898
Report (4-baseline or 7-robustness)	540	0.926	193	0.959
Report (8-robustness)			173	0.960
Report (5-baseline or 9-robustness)	533	0.929	170	0.965
Report (10-robustness)			177	0.983
Secret number (no report)	1,014	1.802	581	2.616
Guess (report=1)	61	1.533	28	1.429
Guess (report=2)	235	2.243	42	2.190
Guess (report=3)	395	3.061	82	3.244
Guess (report=4)	500	4.009	132	4.201
Guess (report=5)	495	4.825	168	5.244
Guess (report=6)			168	6.202
Guess (report=7)			185	7.122
Guess (report=8)			166	8.018
Guess (report=9)			164	8.963
Guess (report=10)			174	9.920
Guess (no report)	1,014	2.283	581	3.419
Guess - secret number (no report)	1,014	0.481	581	0.803

Table 4. Payoff losses for senders and receivers (main sessions, robustness sessions, and high incentives task). In Panel A, the payoff loss is the fraction of the highest expected payoffs that was not achieved for a given decision (relative to the expected payoffs for that decision). In Panel B, the payoff loss is the fraction of the equilibrium payoffs that was not achieved for a given decision (relative to the expected payoffs for that decision). All payoffs have been normalized by the minimum payment that could be achieved with a decision. Expectations are formed treating all secret numbers as equally likely and assuming there is an equal chance of facing players in the other role from any round of the same session. Receiver losses are just for rounds where the secret number is not reported. In Panel C, we calculate the distance from the guess with highest expected payoff divided by the width of the state space (5 or 10).

Panel A: Fraction of highest expected payoffs not earned			
VARIABLES	N	mean	sd
Receiver loss (no feedback & fixed role)	884	0.105	0.146
Receiver loss (no feedback & random role)	1,014	0.132	0.215
Receiver loss (feedback & random role)	675	0.0972	0.177
Receiver loss (robustness sessions)	581	0.106	0.188
Receiver loss (high incentives task)	57	0.0927	0.128
Sender loss (no feedback & fixed role)	2,565	0.0546	0.148
Sender loss (no feedback & random role)	2,700	0.0477	0.137
Sender loss (feedback & random role)	2,025	0.0684	0.169
Sender loss (robustness sessions)	1,890	0.0733	0.154
Sender loss (high incentives task)	285	0.0344	0.110

Panel B: Fraction of equilibrium payoffs not earned			
VARIABLES	N	mean	sd
Receiver loss (no feedback & fixed role)	884	0.222	0.128
Receiver loss (no feedback & random role)	1,014	0.264	0.187
Receiver loss (feedback & random role)	675	0.213	0.163
Receiver loss (robustness sessions)	581	0.199	0.170
Receiver loss (high incentives task)	57	0.211	0.112

Panel C: Fraction of distance from action with highest expected payoffs			
VARIABLES	N	mean	sd
Receiver mistake (no feedback & fixed role)	884	0.157	0.133
Receiver mistake (no feedback & random role)	1,014	0.171	0.170
Receiver mistake (feedback & random role)	675	0.141	0.148
Receiver mistake (robustness sessions)	581	0.157	0.162
Receiver mistake (high incentives task)	57	0.144	0.131

Table 5: Summary of player mistakes relative to a player's self-reported belief and dynamic response to last block of rounds (define every 5 rounds as one block)

Panel A: Sender

	No feedback & fixed role	No feedback & random role	Feedback & random role
# of unique subjects that have acted as sender	57	120	90
Belief of receiver guess if no report (subject-specific, self-reported)	2.273	2.143	1.617
Action consistent with self-reported belief (report if belief \leq draw, not report if belief $>$ draw)	0.875	0.872	0.858
Rounds 1-15	0.856	0.816	0.793
Rounds 16-30	0.876	0.887	0.861
Rounds 31-45	0.891	0.913	0.921
Action consistent with average receiver behavior in last block of rounds (report if average guess given reporting \geq average guess given non-reporting)	0.824	0.805	0.863
Rounds 6-15	0.814	0.773	0.798
Rounds 16-30	0.813	0.814	0.864
Rounds 31-45	0.843	0.817	0.907
Sender loss (fraction of highest expected payoffs not earned given average receiver behavior in last block of rounds)	0.051	0.068	0.065
Rounds 6-15	0.052	0.088	0.085
Rounds 16-30	0.062	0.061	0.074
Rounds 31-45	0.039	0.062	0.043
<i>If draw=1 or 2:</i>			
Action consistent with self-reported belief	0.815	0.831	0.778
Rounds 1-15	0.805	0.762	0.696
Rounds 16-30	0.825	0.853	0.777
Rounds 31-45	0.814	0.885	0.858
Action consistent with avg receiver behavior in last block of rounds	0.665	0.650	0.706
Rounds 6-15	0.679	0.630	0.591
Rounds 16-30	0.663	0.707	0.718
Rounds 31-45	0.659	0.594	0.777

Sender loss	0.085	0.109	0.131
Rounds 6-15	0.083	0.150	0.173
Rounds 16-30	0.098	0.083	0.146
Rounds 31-45	0.072	0.104	0.085

Panel B: Receiver

	Fixed role No feedback	Random role No feedback	Random role Feedback
# of unique subjects that have acted as receiver	57	120	90
<i>Conditional on sender not reporting</i>			
Guess of secret number	2.243	2.283	1.897
Rounds 1-15	2.325	2.464	2.190
Rounds 16-30	2.211	2.251	1.904
Rounds 31-45	2.180	2.088	1.495
Actual secret number	1.734	1.802	1.680
Rounds 1-15	1.811	1.884	1.895
Rounds 16-30	1.781	1.826	1.618
Rounds 31-45	1.590	1.670	1.464
Guess – actual secret number	0.102	0.096	0.043
Rounds 1-15	0.103	0.116	0.059
Rounds 16-30	0.086	0.085	0.057
Rounds 31-45	0.118	0.084	0.006
Mistake: (guess – guess with highest expected payoff given average sender disclosure in last block of rounds)/5	0.159	0.167	0.138
Rounds 6-15	0.167	0.169	0.152
Rounds 16-30	0.146	0.164	0.156
Rounds 31-45	0.169	0.175	0.105
Receiver loss (fraction of highest expected payoffs not earned given average sender disclosure in last block of rounds)	0.107	0.129	0.094
Rounds 6-15	0.113	0.143	0.106
Rounds 16-30	0.093	0.117	0.112
Rounds 31-45	0.118	0.129	0.062

Table 6. Regressions on sender disclosures (main sessions).

Sample Dependent variable	All draws		Only draws of 2		Rounds 6-45 Distance from highest expected payoff (fraction)	
	Report or not		Report or not		(5)	(6)
	(1)	(2)	(3)	(4)		
Dummy=1 if in the first 5 rounds	-0.0442** (0.0173)	-0.0464*** (0.0178)	-0.0652 (0.0422)	-0.0636 (0.0419)		
Round # (1 to 45)	0.00113* (0.000655)	0.00118* (0.000681)	0.00277 (0.00203)	0.00207 (0.00208)	-0.000413* (0.000232)	-0.000417* (0.000232)
Round # * random role * no feedback	0.000686 (0.000876)	0.000329 (0.000899)	0.00378 (0.00258)	0.00225 (0.00271)	1.95e-05 (0.000448)	7.20e-05 (0.000463)
Round # * random role * feedback	0.00180* (0.000939)	0.00167* (0.000955)	0.00706** (0.00278)	0.00651** (0.00300)	-0.000878* (0.000503)	-0.000737 (0.000549)
Dummy=1 if sender belief of receiver guess upon non-report is below the actual draw	0.241*** (0.0360)		0.321*** (0.0561)		-0.0402*** (0.0119)	
Dummy=1 if draw=2	0.219*** (0.0284)	0.335*** (0.0269)			0.00586 (0.0183)	-0.0138 (0.0181)
Dummy=1 if draw=3	0.523*** (0.0351)	0.700*** (0.0246)			-0.0194 (0.0204)	-0.0468** (0.0181)
Dummy=1 if draw=4	0.607*** (0.0365)	0.837*** (0.0219)			-0.0442** (0.0198)	-0.0808*** (0.0162)
Dummy=1 if draw=5	0.613*** (0.0383)	0.839*** (0.0224)			-0.0493** (0.0208)	-0.0847*** (0.0170)
Individual demographics	x	absorbed	x	absorbed	x	absorbed
Session fixed effects	x	absorbed	x	absorbed	x	absorbed
Subject fixed effects		x		x		x
Observations	7,224	7,224	1,477	1,477	5,742	5,742
R-squared	0.512	0.580	0.180	0.629	0.075	0.224

Notes: In parentheses are robust standard errors clustered by subject. *** p<0.01, ** p<0.05, * p<0.1 We define every 5 rounds as one block. Highest expected payoff is based on the distribution of receiver behavior he/she has observed in the last block of the same session. Columns (5) and (6) exclude the first 5 rounds because we need to construct the initial condition from the first 5 rounds. In all regressions, the default is the no feedback & fixed role treatment, and draw=1.

Table 7. Regressions on receiver guesses of non-reported secret numbers (main sessions).

Dependent variable	Receiver guess		Distance from highest expected payoff (fraction)	
	(1)	(2)	(3)	(4)
Dummy=1 if in the first 5 rounds	0.159*** (0.0584)	0.111** (0.0551)		
Round # (1-45)	-0.00358 (0.00247)	-0.00382 (0.00247)	0.000247 (0.000482)	0.000133 (0.000509)
Round # * random role * no feedback	-0.00466 (0.00324)	-0.00564* (0.00324)	-0.000162 (0.000765)	-0.000369 (0.000792)
Round # * random role * feedback	-0.0173*** (0.00374)	-0.0182*** (0.00391)	-0.00197*** (0.000744)	-0.00208*** (0.000774)
Implied average non-reported number given receiver stated beliefs	0.695*** (0.102)		0.120*** (0.0226)	
Individual demographics	x	absorbed	x	absorbed
Session fixed effects	x	absorbed	x	absorbed
Subject fixed effects		x		x
Observations	2,551	2,551	2,204	2,204
R-squared	0.315	0.680	0.204	0.636

Notes: In parentheses are robust standard errors clustered by subject. *** p<0.01, ** p<0.05, * p<0.1. We define every 5 rounds as one block. Highest expected payoff is based on the distribution of sender behavior in the last block of the same session. Columns (3) and (4) exclude the first 5 rounds because we need to construct the initial condition from the first 5 rounds.

Table 8: Regressions on receiver guesses of non-reported secret numbers (main sessions).

Sample	All last guesses		Last guess = 2, 2.5, or 3	
	Dependent variable			
	Guess with no report – last guess with no report			
	(1)	(2)	(3)	(4)
Dummy=1 if in the first 5 rounds	-0.027 (0.052)	-0.014 (0.059)	0.065 (0.063)	0.048 (0.066)
Round # (1-45)	-0.000 (0.001)	-0.000 (0.001)	0.000 (0.002)	-0.002 (0.002)
Round # * random role * no feedback	0.002 (0.002)	0.003 (0.002)	0.003 (0.003)	0.002 (0.004)
Round # * random role * feedback	0.001 (0.003)	-0.000 (0.003)	-0.003 (0.007)	-0.010 (0.008)
Over-guessed last time	-0.191*** (0.054)	-0.249*** (0.070)	-0.072 (0.048)	-0.094* (0.054)
Over-guessed last time * random role * no feedback	-0.129 (0.094)	-0.228* (0.119)	0.028 (0.092)	-0.020 (0.098)
Over-guessed last time * random role * feedback	-0.386*** (0.118)	-0.519*** (0.169)	-0.384*** (0.145)	-0.360* (0.182)
Under-guessed last time	-0.005 (0.050)	0.005 (0.058)	-0.073 (0.059)	-0.095 (0.068)
Under-guessed last time * random role * no feedback	0.103 (0.101)	0.131 (0.118)	0.209 (0.133)	0.200 (0.152)
Under-guessed last time * random role * feedback	0.037 (0.075)	0.015 (0.094)	0.124 (0.157)	0.120 (0.198)
Individual demographics	x	absorbed	x	absorbed
Session fixed effects	x	absorbed	x	absorbed
Subject fixed effects		x		x
Observations	2,287	2,306	1,151	1,154
R-squared	0.078	0.124	0.087	0.386

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 in the “random role” treatments (main sessions).
 Just 5 rounds before and after information intervention.

VARIABLES	No feedback Rounds 41-45		No feedback After information		Feedback Rounds 41-45		Feedback After information	
	N	mean	N	mean	N	mean	N	mean
Report (secret number=1)	66	0.0758	72	0.0278	46	0.0435	44	0.0227
Report (secret number=2)	50	0.620	57	0.474	36	0.667	48	0.792
Report (secret number=3)	75	0.853	47	0.809	48	1	45	0.956
Report (secret number=4)	61	0.967	54	0.963	47	0.979	46	0.957
Report (secret number=5)	48	0.917	70	0.943	48	0.979	42	1
Secret number (no report)	97	1.649	115	1.609	58	1.328	57	1.351
Guess (no report)	97	2.129	115	2.170	58	1.431	57	1.614

Appendix

Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure

Ginger Zhe Jin
University of Maryland and NBER

Michael Luca
Harvard Business School

Daniel Martin
Northwestern Kellogg School of Management

1 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

2 Other Bayes Nash Equilibria

In the disclosure game presented in this paper, nature moves first by selecting the secret number, which creates a game of incomplete information. A standard solution concept for such games is Bayes Nash Equilibrium (BNE). BNE have two defining characteristics: strategies are optimal given beliefs and beliefs are correct (according to Bayes' Rule) *on the equilibrium path*. Because there are no requirements for strategies and beliefs off the equilibrium path, both can be constructed so that otherwise "reasonable" deviations are ruled out. As a consequence, games of incomplete information can end up with a multitude of BNE, some of which seem implausible. For our game, this can produce partial disclosure BNE.

Definition: A partial disclosure Bayes Nash Equilibrium (BNE) of our game is a BNE in which senders of multiple types put a positive probability on not reporting.

For instance, when the set of receiver actions is sufficiently rich, there can exist a BNE of our disclosure game in which senders of all types do not report and receivers take an 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 closest to the bottom of the state space is taken when the sender reports, regardless of what the sender reports. Clearly, guessing 1 when the sender has reported 5 is not optimal, but because senders never report in this equilibrium, that node of the extensive form game is never reached, so the receiver's payoffs are not improved by deviating.

To show that the data in our experiment are not consistent with *any* partial disclosure BNE, we show that there is a condition that is inconsistent with all such equilibria. This condition is clearly evident in our experimental data, so we can conclude that we do not observe behavior that is consistent with a partial disclosure BNE.

Proposition: In every partial disclosure BNE of our game, it cannot be that all secret numbers above the minimum are reported with a positive probability.

This proposition is simple to prove. By definition, every partial disclosure BNE has at least two types that choose non-reporting with a positive probability. Thus, the Proposition is equivalent to saying it is not possible (in a BNE) to simultaneously have more than one type not reporting with positive probability (partial disclosure) and all types reporting with positive probability.

Because on-path beliefs must be correct in a BNE, the guess of the non-reported number must be below the highest type. By contradiction, assume that all numbers above the

minimum are reported with a positive probability. In that case, the highest type that puts a positive probability on not reporting always has a profitable deviation to reporting, since the receiver will be guess the secret number in that case. But this cannot be true in a BNE.

In our data, we observe a positive probability of reporting in every secret number above the minimum, hence according to the Proposition, it cannot be a partial disclosure BNE.

3 NYU Robustness Sessions

We ran additional sessions of the “no feedback & random role” treatment with five secret numbers (as in the main sessions) in the Center for Experimental Social Science (CESS) laboratory at New York University (NYU). These sessions were excluded from our main sessions because beliefs were not elicited from subjects in these sessions. Instead, we use these sessions to perform additional robustness checks for our results that do not require data on beliefs. We also use them to examine behavior in other additional tasks.

3.1 Experimental Design: Other Additional Tasks

In addition to the “aggregate feedback” additional task, we ran five other additional tasks at NYU that do not appear in our main sessions. In the first, which we call the “risk” task, subjects completed the well-known measure of risk aversion introduced by Holt and Laury (2002). 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.

The aim of this task was to see whether subject choices were related to the risk preferences of subjects. Risk preferences can impact receiver guesses when there is uncertainty about the underlying state, and risk aversion can push guesses of non-reported secret numbers higher when the distribution of non-reported secret numbers is skewed towards lower numbers because higher guesses produce lower variation in payments. For instance, given the overall reporting rates for the “fixed role & no feedback” treatment reported in Table 2, the unconstrained optimal guess for a risk neutral agent would be 1.6156, and for a risk averse agent with the preferences $U(x)=x^{.75}$, the optimal guess would be 1.6725. Note that the predicted difference is small, even for substantial changes in risk preferences, so we might not expect to see a strong relationship between receiver guesses and risk preferences.

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 payoff implications for others. By comparing these choices with previous choices, we can determine whether sender and receiver choices were impacted by social concerns related to the payoffs of the subject they were paired with.¹ Niederle and Vesterlund (2007) use a similar approach to separate preferences for competition from social preferences. Note that this task is identical to the “high incentives” task, but with a normal payoff rate. As in the “high incentives” task, guesses in this task are potentially impacted both by changes in the payoff implications for senders and by the fact that behavior may be changing over rounds because of learning.²

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. However, guesses in this task are also potentially impacted by the fact that behavior may be changing over 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.

In fifth task, which we call the “average reports” task, subjects were shown the average *reported* secret number from all subjects in that session from the first 45 rounds and then completed the same steps as in the “aggregate feedback” task. Because the number of

¹ The various ways in which social considerations could potentially impact receiver guesses in our game is discussed in this appendix.

² The latter produces several possible confounds. For example, subjects may have learned to play differently, or subjects may have changed their beliefs about what other subjects have learned. We thank an anonymous referee for pointing this out.

rounds in which the secret number was reported was not provided, there was not enough information for subjects to fully pin down the average non-reported secret number. For instance, the average reported secret number would be 4 if the secret number was reported in just one round where the secret number was 4 (leaving the average non-reported secret number near 3), and the average reported secret number would also be 4 if all senders with secret numbers of 3, 4, and 5 reported (leaving the average non-reported secret number near 1.5). However, by placing additional assumptions on the actions of senders, more information can be gleaned from the average reported secret number. For instance, by assuming monotonic reporting rules and the same cut-off for all senders, the average non-reported secret number can be fully pinned down from the average reported secret number.

3.2 Results: Robustness

In our sessions at NYU, we also used a show-up fee of \$5, and on average subjects earned \$25.25. Table A1 shows the summary statistics for the NYU sessions. All 212 subjects at NYU were assigned to the “no feedback & random role” treatment. NYU subjects were more likely to self-report as undergraduates, female, and non-native English speakers than our HBS subjects.

Table A1. Summary statistics sessions at NYU. Observation is per subject. Value is missing if demographic information not provided by the subject.

VARIABLES	N	mean
Number of subjects in the session	212	13.40
Feedback provided (dummy)	212	0
Random role (dummy)	212	1
Undergraduate (dummy)	212	0.858
Male (dummy)	212	0.368
Native English speaker (dummy)	212	0.698
Friend in the session (dummy)	212	0.0943

Table A2 compares the actions of senders and receivers across schools. At NYU, there is less disclosure of less favorable draws and more disclosure of more favorable draws, which produces a lower average secret number when senders do not report. Also, at NYU receivers are more pessimistic about non-reported secret numbers. However, the extent to which receivers overestimate non-reported secret numbers is very similar between schools (0.481 and 0.448) and is not statistically significant at the 10% level using a 2-sided t-test.

Table A2. Summary of player actions for “no feedback & random role” sessions at HBS and NYU. Note: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

VARIABLES	HBS		NYU		p-value (=)
	N	mean	N	mean	
Report (secret number=1)	568	0.107	929	0.0614	0.0013***
Report (secret number=2)	552	0.426	986	0.398	0.2813
Report (secret number=3)	507	0.779	973	0.860	0.0001***
Report (secret number=4)	540	0.926	960	0.977	0.0000***
Report (secret number=5)	533	0.929	922	0.972	0.0001***
Secret number (no report)	1,014	1.802	1,650	1.628	0.0000***
Guess (report=1)	61	1.533	57	1.219	0.0551*
Guess (report=2)	235	2.243	390	2.097	0.0005***
Guess (report=3)	395	3.061	836	3.039	0.2622
Guess (report=4)	500	4.009	937	4.012	0.8409
Guess (report=5)	495	4.825	895	4.960	0.0000***
Guess (no report)	1,014	2.283	1,648	2.076	0.0000***
Guess - secret number (no report)	1,014	0.481	1,648	0.448	0.5060

3.3 Results: Other Additional Tasks

In the NYU sessions, 38 subjects 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 aggregate feedback task.³ Here we examine results for the risk task and the other task, as they were designed specifically to examine other forces besides belief biases that could explain receiver over-guessing. Results for the other tasks are available in Jin, Luca, and Martin (2015).

We first look at the 38 subjects who completed the risk additional task. 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

³ Unlike the additional tasks completed in our main sessions, the addition tasks completed in our NYU sessions are potentially under-powered because even though there were 212 subjects in total, each subject only completed one additional task and there were six possible additional tasks.

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. 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 positive, but is small (0.031) and not significant ($p=0.261$).⁴

Second, we explore the role of social preferences in guessing of non-reported secret numbers. For evidence of this, we examine 26 NYU subjects who completed the other additional task. As mentioned previously, these subjects guessed the secret number from an earlier round, but now without payoff 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, $p=0.1866$).⁵

The increase in guesses after social considerations are minimized provides suggestive evidence of a punishment motive towards those who do not disclose. Instead of providing a force pushing away from equilibrium (higher guesses for non-reported secret numbers), social considerations appear to be pushing behavior towards equilibrium (lower guesses for non-reported secret numbers).

⁴ However, there is just a single data point per subject, so it should be note that this analysis may be underpowered.

⁵ Once again, there is just a single data point per subject, so this analysis may be underpowered.

4 Social Preferences

Social preferences can impact receiver behavior differently depending on whether or not the sender discloses the secret number, and when the sender discloses, the realization of the secret number. The realization of the state matters because it impacts the distribution of payoffs and individuals have been shown to hold preferences over the distribution of payoffs.

First, when senders *disclose* the secret number, social preferences could lead receivers to over-guess very low secret numbers. 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 instance, standard models of fairness would say that some individuals could feel “bad” about accurately guessing the revealed state when the state is 1 (payoffs: -29, 110) due to feelings of guilt. As shown in Table 2, there is suggestive evidence for this effect, but the size is small for draws of 2, so it does not impact the strategic incentives for unraveling.

Second, when senders *do not disclose* the secret number, the same forces of guilt may be at work. However, because receivers are now uncertain of the state, we have to account for the interaction between risk and social preferences. We know of no model of fairness under risk that would suggest that the impact of social motivations would get stronger with risk. As a consequence, we would expect the impact of guilt to be no greater with non-disclosure than with disclosure. Thus, given the small size of the possible impact of social preferences in the case of disclosure, it seems unlikely that social preferences are driving a substantial part of the over-guessing when senders do not disclose.

On top of this, if there is a social norm of disclosing, then receivers might wish to punish non-disclosure, even if the chances of re-matching are low. We find evidence of just such behavior in the “other” additional task conducted in our NYU sessions. These punishments should further dampen receiver guessing with non-disclosure.

5 References

- 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., Luca, M., & Martin, D. (2015). Is no news (perceived as) bad news? An experimental investigation of information disclosure. NBER Working Paper 21099.
- Niederle, M., & Vesterlund, L. (2007). Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics*, 122(3), 1067–1101.