

NBER WORKING PAPER SERIES

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

Ginger Zhe Jin
Michael Luca
Daniel Martin

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

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

We would like to thank seminar audiences at PSE, Alicante, Wisconsin, Olin, Michigan, Pittsburgh, Rotman, Kellogg MEDS, Queen Mary, UCLA, HBS, and Bocconi for valuable comments. We thank Patrick Rooney for excellent research assistance. All errors are ours. All rights reserved. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Ginger Zhe Jin, Michael Luca, and Daniel Martin. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

ABSTRACT

A central prediction of information economics is that market forces can lead businesses to voluntarily provide information about the quality of their products, yet little voluntary disclosure is observed in the field. In this paper, we demonstrate that the inconsistency between theory and reality is driven by a fundamental failure in consumer inferences when sellers withhold information. Using a series of laboratory experiments, we implement a simple disclosure game in which senders can verifiably report quality to receivers. We find that senders disclose less often than equilibrium would predict. Receivers are not sufficiently skeptical about undisclosed information – they underestimate the extent to which no news is bad news. Senders generally take advantage of receiver mistakes. We find that providing disclosure rates by quality score helps to improve receiver inferences.

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

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

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

1 Introduction

Across a wide range of settings, sellers have private information about the quality of the goods and services they sell. Restaurant owners know the results from their hygiene inspections. Car manufacturers know the gas mileage of their cars. Salad dressing companies know how many calories their dressings contain. Truth-in-advertising laws stipulate that companies cannot provide misleading or incorrect information to customers. In other words, the information they provide must be *verifiable*. However, businesses can often decide whether to provide such information to buyers. When a business chooses not to disclose information, customers must then infer whether no news is bad news or good news.²

Theories of voluntary disclosure, dating back to Viscusi (1978), Grossman and Hart (1980), Grossman (1981), and Milgrom (1981) suggest that market forces can drive firms to voluntarily and completely reveal information about their quality when such information is verifiable and the costs of verification and disclosure are low. The mechanism behind these results is simple. Consumers treat all non-reporting firms the same, so the highest quality non-reporting firms 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 when information is not disclosed.

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 unraveling result suggests that the same benefits 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

² See DellaVigna and Gentzkow (2010) for a general review of persuasive communication.

complete (see Mathios 2000, Jin 2005, Fung et al. 2007, and Luca and Smith 2015 for specific examples). 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: it may be costly for sellers to collect and disclose information to the public, consumers may already know information from other channels, or it may be difficult to disclose the information in a format that is comprehensible to consumers. The second strand focuses on a seller's strategic incentives: sellers may choose not to obtain data on product quality in order to avoid future demand for disclosure (Matthews and Postlewaite 1985), a desire for product differentiation may dominate disclosure incentives (Board 2009), and when the quality information is coarse, sellers of best quality may use non-disclosure as a counter-signal to distinguish themselves from eager-to-disclose medium-quality sellers (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 show that information fails to unravel even after we strip away all the above-mentioned factors in a well-controlled laboratory setting. In fact, our results suggest that the lack of voluntary disclosure we observe in the field may be explained by a more fundamental reason. A crucial element for unraveling to occur is that consumers need to correctly infer the quality of non-disclosing firms – to correctly identify when “no news is bad news” – and our experiment indicates that consumers may underestimate the extent to which no news is bad news. The resulting opportunity to

mislead consumers can incentivize firms to withhold information more often than is predicted by equilibrium.³

In the setting we study, 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. This is in contrast with existing experiments on strategic information transmission (Cai and Wang 2006, Wang, Spezio, and Camerer 2011), where senders can engage in "cheap talk".⁴

After the sender decides whether or not to disclose their private information, the receiver must guess the state.⁵ If the sender has revealed the state through disclosure, then this task is trivial – the receiver knows the true state with certainty. If the sender has not revealed the state, then the receiver must infer it based only on the sender's decision not to disclose. However, the receiver also knows the distribution of states and that the sender has private information. The sender and receiver do not have aligned interests: the sender has higher earnings when the receiver guesses that the state is higher (guesses and states are numeric values), and the receiver has higher earnings when his or her guess is closer to the true state.

The unique sequential equilibrium of this game can be found with a straightforward application of the unraveling arguments mentioned previously. In equilibrium, senders always reveal their information (unless the state takes the lowest possible value, in

³ Brown, Camerer, and Lovo (2012) find that consumers fail to appreciate the extent to which a "cold opening" is bad news about the quality of a film.

⁴ Only a handful of papers have studied verifiable information disclosure in the lab. Most of these studies (for example, Forsythe et al. 1989, King and Wallin 1991, and Dickhaut et al. 2003), are motivated by disclosure in asset markets, so their experimental designs are substantially different from ours. For example, they have receivers compete with each other through an auction mechanism, which introduces room for other biases to drive receiver choices. See section 3 for a comparison of our experiment with other related experiments.

⁵ Guessing the state is analogous to deciding how many units to purchase, as in Milgrom (1981).

which case they are indifferent between revealing and not), and receivers correctly guess that the state takes the lowest possible value when senders do not reveal this information.

When we implement this game in the laboratory, we find widespread failures of unraveling, despite the simplicity of the strategic interaction. Senders do not fully disclose the state, and receivers are not fully skeptical about non-disclosure. To ensure that participants understand the game, we allow participants to play for 45 rounds (being randomly re-matched with a different, anonymous partner in each round) and to play both roles (being randomly assigned a role in each round), but the failures of unraveling we observe are persistent.

We find that most receivers consistently guess higher than predicted by equilibrium and that many receivers guess far higher than the average non-disclosed state. We present evidence that these high guesses result from poor inferences on the part of receivers and provide robustness tests to rule out alternative explanations, such as risk aversion, social preferences, and random choice errors.

These overestimates cause a breakdown in the logic of unraveling because senders with relatively favorable information do not have an incentive to disclose more often. We investigate several ways to increase sender incentives to disclose by improving receiver inferences. Clearly, one option is simply to tell receivers the sender's true type after each round. While we implement this experimentally and find it to be an effective method, it is not a realistic policy intervention for most settings of interest. For example, it would be a strange policy initiative to tell all consumers the actual hygiene grade of the restaurant after they eat their meal.

We implement two other interventions that do correspond to feasible policy initiatives. First, we tell receivers the average state when senders chose to disclose (as opposed to when they chose to withhold). This contains some information about the strategies of senders; however, we find that it does not affect receiver beliefs or the rates of disclosure. Second, we tell receivers the fraction of senders who chose to report

and not report for each state. We find that this intervention leads to improved guesses on the part of receivers and higher rates of disclosure on the part of senders.

Overall, our lab experiments demonstrate that allowing sellers to verifiably and costlessly disclose information about product quality is not sufficient for unraveling because customers do not sufficiently interpret the negative signal of non-disclosure. As a result, a critical question is how to educate consumers so that they will interpret no news as bad news. Our experiments have shown that a simple summary of disclosure rates by quality score may go a long way towards overcoming incorrect buyer beliefs. Kessler and Roth (2012) argue that laboratory experiments can be used as a starting point for policy interventions, and we hope our laboratory experiments can form the starting point for new policies aimed at increasing voluntary disclosure.

The rest of the paper is organized as follows. Section 2 describes the disclosure game, section 3 provides details on how we implement this game with laboratory experiments, section 4 reports the experimental results, and section 5 presents an additional experiment used to test the robustness of our results, and section 6 concludes with policy implications.

2 The Disclosure Game

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

The sender has two possible actions, and the receiver is aware that these are the only two actions available to the sender. The sender can either report the state to the receiver or make no report. This report must be truthful and cannot be vague. Thus, the set of actions M available to a sender of type s is just $M(s)=\{s,null\}$.

Regardless of whether or not they receive a report from the sender, the receiver

takes an action a from a finite space A , which is also a subset of the real numbers and contains S . We interpret this action as guessing the type of the sender. We could have considered a setting where the receiver's action is to choose a quantity to purchase (as in Milgrom 1981), but this would have added more complication to the game.

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

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

There are other Bayesian Nash equilibria of this game, but they require strategies to contain non-rational behavior off of the equilibrium path. For instance, there can exist a Nash equilibrium in which sellers never report and receivers take the action that is as close to the average realization of the state space as possible. This is supported by a receiver strategy in which the action equal to the minimum element of the state space is taken if the sender does report. However, it would not be optimal for the receiver to take this action if the state was reported and it was not the minimum element of the state space. Importantly, we do not observe behavior consistent with any of these other Bayesian Nash equilibria.

Finally, we think it would be interesting to consider disclosure games where there is more than one sender or receiver, where there is more than one period, where the sender is potentially uninformed about the true state, where the sender also has access to vague or untruthful messages, or where the sender's preferences are not monotonic in the receiver's action. However, our aim is to study the simplest possible setting of verifiable disclosure where unraveling is predicted to generate voluntary disclosure, so we leave these extensions for future work.

3 Experimental Design

In our experiment, subjects completed 45 rounds and then, depending on the session, one of seven possible additional tasks. Subjects were told at the beginning of the experiment that they would complete an additional task, but were given no details about the task. See the appendix for the full set of instructions given before start of the experiment.

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 ECU). At the end of the experiment, each subject's ECU balance was rounded up to the nearest non-negative multiple of 200 and converted into U.S. dollars at a rate of 200 to 1.

3.1 In Each Round

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

In each round and for each pairing, one subject was randomly assigned to be the sender, and the other subject was randomly assigned to be the receiver. Each was

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

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

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

Once all senders had 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, and that if the area for messages on the receiver’s screen was left blank, it was because the sender did not report the secret number.

Below the area for messages, receivers were asked to guess the secret number, and these guesses could be any half unit between 1 and 5. Thus, the set of actions is $A=\{1,1.5,2,2.5,3,3.5,4,4.5,5\}$. The actions of receivers were limited to half unit increments so that payoffs could be represented in a table. There was also no time limit for receiver decisions either.

Receiver payoffs in each round were $ECU_R = 110 - 20|S - A|^{1.4}$, where S is the secret number and A is the receiver’s guess. These payoffs are such that a risk neutral receiver would guess closest to their expected value of the secret number. The exact sender payoffs in each round were $ECU_S = 110 - 20|5 - A|^{1.4}$. These payoffs are independent of the secret number and monotonically increasing with receiver actions,

because guesses could not be higher than 5. Because there was a small number of states and actions, the payoffs could be shown in a table, so that subjects did not need to know or interpret these functional forms.

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

As possible extensions of this experiment, it might be interesting to consider different prior probabilities of each state, a larger state space, or a finer grid of actions. However, our goal was to keep the design of our experiment as simple as possible, so as to isolate the strategic tension of interest.

3.2 Between and Across Periods

We used the payoff functions described above because they produce the desired strategic tensions, but also because they produce reasonable final payoffs. We paid subjects cumulatively for their decisions in every round, which could result in deliberate variation in play (often called a “portfolio” strategy). However, there is little evidence of such behavior in our experiments or in experiments that use the same payoff functions and also pay cumulatively (such as Wang, Spezio, and Camerer 2010).

Also, we did not provide feedback after each round about the actual secret number in that round, about the receiver’s guess in that round, or about the payoffs in that round. Excluding such feedback is an important and intentional aspect of our experimental design. We did this to mirror many of the settings in which voluntary and verifiable disclosure is studied in the field. In some cases, consumers only have limited experience in a market, so they do not have the chance to acquire feedback. In others, consumers have long experience but get little or very noisy feedback about non-reported product quality. For example, in the case of salad dressing nutritional labels, it would be difficult for a consumer to know that they have eaten a salad dressing with inferior nutritional characteristics.

3.3 Related 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).

However, there is one substantial difference in our experimental design: in our experiment, the sender's messages must be truthful. This is why our experiment is a test of verifiable disclosure, and their experiments are tests of cheap talk.

There are only a limited number of experiments that test verifiable disclosure, and as discussed in the introduction, there are important differences between the designs of these experiments and ours.

Three of these papers (Forsythe, Isaac, and Palfrey 1989, King and Wallin 1991, and Dickhaut, Ledyard, Mukherji, and Sapa 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 assets 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 Sapa (2003) complement these findings by also showing what happens when there is the possibility that senders may not be informed about the asset's quality. The latter goes beyond the first by considering both partially informed senders and partially informative messages.

As mentioned previously, these experiments represent a valuable test of disclosure in asset markets, but they are less applicable to our settings of interest. Also, the use of auctions introduces the room for other biases to drive disclosure decisions, particularly since these experiments used first-price auctions.

In addition, Forsythe, Lundholm, and Rietz (1999) compare disclosure to cheap talk in reducing adverse selection. 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. They too find that reports converge to full unraveling.

Concurrent to our study are three new papers that use experiments to study verifiable disclosure. Bhattacharya, Kang, and Wilson (2015) aim to more closely compare disclosure to cheap talk, and they allow senders to disclose an interval of states. 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 our experiments, the receiver is a computer who uses an automated strategy, so there is no room for inference problems. Hagenbach and Perez-Richet (2015) investigate a simple verifiable disclosure game where sender payoffs are not necessarily monotonic, but feature incentives for senders to masquerade as another type.

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 our experiment, in Sanchez-Pages and Vorsatz (2009) non-disclosure carries a cost. Even with this cost, some senders choose not to disclose. Serra-Garcia, van Damme, and Potters (2011) find that intermediate senders sometimes use vague messages, which receivers do not make correct inferences about. Agranov and Schotter (2012) also study the use of vague language but focus on the vagueness possible with human language.

One substantial difference between our design and the design of the existing voluntary disclosure experiments is that we do not provide full feedback after each round. As we will describe in more detail later, we conducted an additional experiment in which full feedback was provided, and consistent with much of the existing literature, we found convergence towards full unraveling. However, even in markets with many repeat purchases, we do not feel that full feedback is either (1) a realistic description of the feedback consumers receive or (2) a realistic policy prescription for most settings of voluntary and verifiable disclosure.

3.4 Additional Tasks

After completing 45 rounds, each subject faced one of seven additional tasks. In the first additional task, which we call the “Distribution” task, subjects were asked to guess the rate at which senders reported each secret number in the initial 45 rounds of the experiment. The aim of this task was to assess whether subject beliefs about sender strategies were correct. The guesses in this task were not incentivized, which introduces the potential for extra noise in the responses.

The second additional task, which we call the “Self” task, is one in which subjects played once more in the role of sender and in the role of receiver, but this time against *their own* decisions from past rounds. When in the role of sender, subjects were told that if they reported, their computer opponent would guess the secret number, and if they did not report, their computer opponent would match the guess they had made in a past round in which their opponent did not report. When in the role of receiver, subjects were asked to guess the secret number from a previous round in which they did not report the secret number. The payoffs from this task were added to the ECU earned in the first 45 rounds.

It seems possible, but unlikely, that subjects could remember all of their past decisions, so subjects would need to remember their own past strategy. 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.

The third, which we call the “Holt-Laury” 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.

The aim of this task was to see whether sender and receiver choices were related to the risk preferences of subjects. At the end of the experiment, one choice was randomly selected, and any earnings from the realization of that lottery were added to the show-up fee and earnings from the first 45 rounds.

We call the fourth additional task the “Past” task, because after 45 rounds subjects played once more in the role of sender and in the role of receiver, this time against opponents’ decisions from past rounds. When in the role of sender, subjects were told that if they reported, their computer opponent would guess the secret number, and if they did not report, their computer opponent would match the guess of one of their guesses of a receiver from a past round in which they did not report. When in the role of receiver, subjects were asked to guess the secret number from a previous round in which their opponent did not report the secret number. The payoffs from this task were added to the ECU earned in the first 45 rounds.

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 any social preferences. Niederle and Vesterlund (2007) use a similar approach to separate preferences for competition from social preferences.

In the fifth, 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) will report the secret number 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. 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 good assumption for small firms, but not necessarily large firms.

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

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

We call the final additional task the “Consumer Reports” task because the information that subjects were shown was in the style of the popular publication of the same name. Subjects were shown the number of times that each secret number was reported and not reported for all subjects from the first 45 rounds. This provided enough information to determine both the average reported secret number and the average non-reported secret number.

4 Experimental Results

Our primary study was conducted in the Center for Experimental Social Science (CESS) laboratory at New York University. In this laboratory, subjects are separated with dividers, and each subject is provided with a personal computer terminal. Our experiment was run using the z-Tree software package (Fischbacher 2007).

In this study, we observed 212 subjects complete a total of 9,540 rounds. Over 16 sessions, the median and mode session size was 14 subjects. All subjects were students at New York University. At CESS we used a show-up fee of \$5, and on average subjects earned \$25.40. The minimum payment was \$18 and the maximum payment was \$30.85.

As we show in section 4.2, there appears to be learning in the first 5 rounds. Because feedback is limited, it is likely that subjects are learning about the structure of the game, such as internalizing the payoffs in each role and confirming that senders cannot misrepresent the state. Because we want to study how subjects play once they know the structure of the game, we drop the first 5 rounds of each session in the analyses that follow. However, our results are robust to including these 5 rounds.

4.1 Senders Partially Disclose

Looking first at sender behavior, we find that while disclosure is partial, it is far from random. Figure 1 displays the average rate of reporting for each possible secret number. A secret number of 1 is reported 5.7% of the time, a secret number of 3 is reported 88.6% of the time, and secret numbers of 4 or 5 are reported over 97.7% of the time. Most of the action occurs for a secret number of 2, which is reported in 40.8% of the time it is realized.⁶

At these reporting rates, if each secret number was realized with exactly equal probability, the average non-reported secret number would be 1.569. The true average non-reported secret number is 1.584, with a standard deviation of 0.770. Looking across sessions, the standard deviation in the average non-reported secret number is 0.150.

Most of the aggregate variation in reporting a secret number of 2 is due to heterogeneity between individuals, not variability within each individual. 42.9% of subjects never report a secret number of 2, and 25% of subjects always report a secret number of 2. The average standard deviation in the reporting rate within each individual is 0.168, but the median is 0. Given the stability in individual sender behavior over the course of the experiment, it follows that aggregate sender behavior is stable over the

⁶ Looking at the individual level, just 2.4% of subjects always reported the secret number. This is far lower than some estimates in the experimental literature for the fraction of “honest” types. Such a discrepancy could be due to how reporting is framed in our experiment or in differences between how subjects view “lies of omission” and active lying. Exploring these issues is a potentially interesting avenue for future work.

course of the experiment. This may not be surprising given the lack of feedback between rounds.

As mentioned previously, failing to report a secret number of 2, 3, 4, or 5 goes against the predictions of the unique sequential equilibrium. Looking just at rounds where the realized secret number was one of these four numbers, the average percentage of non-equilibrium choices for each subject when in the role of sender is 19.6%, with a standard deviation of 17.4%. The median percentage is 16.7%, the 25th and 75th percentiles are 5.0% and 31.3% respectively, which shows that while some subjects made many more non-equilibrium choices, most subjects made some non-equilibrium choices when in the role of sender.

However, given that the average guess for a non-reported secret number is 2.022, it is a (risk neutral) sender's best response not report a secret number of 1 and to report all secret numbers above 2. Given that the 95% confidence interval for the average guess for a non-reported secret number is 1.978 to 2.066, we will consider both reporting and not reporting a best response when the secret number is 2.

If we look just at secret numbers of 1, 3, 4, and 5, the average percentage of choices that are not a best response for each subject when in the role of sender is 5.3%, with a standard deviation of 11.1%. The median percentage is 0%, the 25th and 75th percentiles are 0% and 6.1% respectively, which shows that most subjects best respond when in the role of sender. This stands in strong contrast to the widespread departures from equilibrium.

4.2 Receivers Guess Too High

Turning to receiver behavior, we find systematic deviations from the predictions of equilibrium when the secret number is not reported.⁷ However, unlike sender behavior, receiver choices are often not a best response.

⁷ Subjects had no problem guessing correctly when the secret number was actually reported by senders, as 92.2% of guesses were identical to the reported number. 57.3% of the remaining guesses were just .5 higher, which appear to be small "rewards" for

Figure 3 shows the aggregate distribution of guesses when the secret number was not reported. The largest mass point is for a guess of 1.5, which is guessed in 25.4% of the time. However, 1 is guessed 19.7% of the time, 2 is guessed 21.5% of the time, and 3 is guessed 18.1% of the time.

As mentioned previously, the average guess across rounds is 2.022, and this number is fairly stable across the course of the experiment. Figure 5 shows the average guess for each block of 5 rounds (after the first 5). The horizontal line is the average guess across rounds (after the first 5). There appears to be a slight decrease in the average guess across rounds, which is confirmed by regressing the guess on the round number (controlling for subject fixed effects and using robust standard errors).⁸ However, the effect size is very small (-0.002) and not significant at a 10% level ($p=0.132$).⁹

As mentioned previously, the prediction of equilibrium is that receivers will always guess the secret number is 1 when it is not reported. However, 81.3% of choices do not correspond to equilibrium. At the individual level, 91.5% of subjects make at least one non-equilibrium choice, and 60.4% of subjects never make the equilibrium choice.

The average non-reported secret number is 1.584, with a 95% confidence interval of 1.544 to 1.624. Given that the average non-reported secret number is closest to 1.5, the best response for (risk neutral) receivers is to guess 1.5, which is above the equilibrium prediction of 1. While 25.4% of guesses correspond to this best response, only 8.5% of subjects always best respond, and 55.2% never best respond.

reporting, and 10% were guesses under 5 for reports of 5, which appear to be a “punishments” for getting a high secret number.

⁸ If we regress guess onto both the number of rounds as sender and the number of rounds a receiver (controlling for subject fixed effects and using robust standard errors), the coefficient on rounds as sender is more negative (-0.010 to -0.056) and has a smaller p-value (0.076 to 0.289). However, both effect sizes are small and neither is significant at a 5% level.

⁹ On the other hand, there is stronger evidence that guesses drop over the first 5 rounds. Restricted to the first 5 rounds, if we regress guess onto period (with subject fixed effects and robust standard errors), the coefficient for period is significant at a 5% level ($p=0.045$). However, this coefficient is still relatively small (-0.052).

However, because of the shape of the receiver's payoff function, the lost earnings from not best responding differ substantially by the guess made. Given the actual reporting rates for each secret number, if each secret number occurred with exactly equal probability, the returns to guessing 1.5 when the secret number is not reported are 98.5 ECU. The lost earnings associated with a guess of 1 would be 1.8, with a guess of 2 would be 2.75, with a guess of 3 would be 25.8, with a guess of 4 would be 60.6, and with a guess of 5 would be 102.7.

4.3 Replication Study

We conducted a replication study at the Computer Lab for Experimental Research (CLER) facility at the Harvard Business School. Once again the experiment was run using the z-Tree software package (Fischbacher 2007). At CLER we also used a show up fee of \$5 and kept the same conversion rate of ECU to dollars.

In this study, we observed 120 subjects complete a total of 5,400 rounds. Over 9 sessions, the median and mode session size was 14 subjects (as before). Here also, we drop the first 5 rounds of each session. This still leaves 40 rounds per subject, for a total of 4,800 rounds.

Unlike the subject pool at CESS, the subject pool at CLER includes both students and non-students. The only requirement for participation was not to be a Harvard University employee. Of the 120 subjects who completed our experiment, just 55% were undergraduates in their first four years of studies.

Figure 2 displays the average rate of reporting for each possible secret number in our replication study. In comparison to Figure 1, the rate of reporting for secret numbers of 1 appears slightly higher in the replication study, and the rate of report for secret numbers of 5 appears slightly lower. However, a Pearson's chi-square test cannot reject the null hypothesis that the true reporting rates are the same in both populations ($p=0.342$).

Figure 4 shows the aggregate distribution of guesses when the secret number was not reported in our replication study. In comparison to Figure 3, the fraction of rounds

where subjects guess 1.5 appears smaller and the fraction of rounds where subjects guess 5 appears slightly larger. However, once again the results of the replication study are not significantly different from the results of the primary study. A Kolmogorov-Smirnov test cannot reject the null hypothesis that the true distributions of guesses are the same in both populations ($p=0.361$).

4.4 Exploring Receiver Over-Guessing

Not only do receivers fail to best respond, they more often guess higher than the best response – sometimes far higher. As a result of this “over-guessing”, when the secret number is 2, senders do not have an incentive to separate through disclosure. This exposes a fundamental breakdown in the mechanics of unraveling.

Thus, we spend most of the remainder of the paper exploring the nature and potential causes of receiver over-guessing. Our leading explanation is that receivers have incorrect beliefs that stem from bad inferences about non-disclosure. We also consider and reject several alternative explanations: stochastic errors, choice heuristics, risk aversion, and social preferences.

4.4.1 Bad Inferences about Non-Reports

Our most convincing evidence that bad inferences about non-reports drive over-guessing is that receiver guesses about non-disclosed secret numbers improve substantially after being provided additional information about reporting. This evidence is discussed in more detail when we report the results of our “Consumer Reports” intervention.

For policymakers, it may be sufficient to know that bad inferences by consumers can produce incomplete disclosure. However, to design effective information interventions, it may be helpful to know the source of these mistakes.

One possible reason for bad inferences about non-disclosure is that receivers have incorrect beliefs about the reporting strategies of senders. Another possible explanation is that receivers have trouble conditioning on actions, so even though they have correct

beliefs about the reporting strategies of senders, they fail to correctly account for the informational content of actions.

We use our first additional task (the “Distribution” task) to help separate these two possible explanations, and we find evidence in support of conditioning failures. Our second additional task (the “Self” task) provides additional supporting evidence for this conclusion. We then show how a leading theory of conditioning failures, Cursed Equilibrium (Eyster and Rabin 2005), can explain the behavior of most subjects in our experiment.

All 120 subjects in the replication study completed the “Distribution” task, in which subjects guess the reporting rate for each secret number after the initial 45 rounds were complete. By looking directly at these responses, we can get a sense for whether subjects have incorrect beliefs about sender strategies, which is the foundation of behavioral approaches such as “Level-k” theory. For example, if a receiver believes that 20% of senders with secret number 1 will report, but actually only 5% of such senders report, this receiver does not have correct beliefs about sender strategies.

Also, if subjects correctly use Bayes’ rule, then combining their beliefs about sender strategies with the probability of each secret number implies a belief about non-reported secret numbers. By looking at the differences between these “implied beliefs” and what subjects actually guessed when secret numbers were not reported, we can get a sense for whether subjects are failing to correctly condition on actions. For example, if a receiver believes that the sender reporting rate is 20% for a secret number of 1, 40% for 2, 80% for 3, and 100% for 4 and 5, then she should guess the secret number conditional on non-reporting as 1.625 because $[(100\%-20\%)*1+(100\%-40\%)*2+(100\%-80\%)*3]/[(100\%-20\%)+(100\%-40\%)+(100\%-80\%)]=1.625$. If when faced non-disclosure her actual guess is far from 1.625, then she has a problem translating her beliefs about sender strategies into a guess conditional on non-reporting.

Figure 6 shows the frequency of guesses for the reporting rates, where a larger bubble represents more subjects guessing closer to that point. Looking across subjects, the median guess of the reporting rate is 3.5% for a secret number of 1, 20% for 2, 50%

for 3, 95% for 4, and 100% for 5 (actual rates were 10.7%, 42.6%, 77.9%, 92.6%, and 92.8% respectively). At the individual level, the median difference between a subject's guess of a reporting rate and the actual reporting rate is 10.7 percentage points. So while there is evidence of incorrect beliefs about sender strategies, it appears that most subjects understand the broad features of actual reporting rates.

On the other hand, we see clear evidence of conditioning failures. As shown in Figure 7, the distribution of mode guesses from the first 45 rounds is shifted to the right of the implied beliefs about the average non-disclosed number. In fact, implied beliefs are on average 0.352 lower than mode guesses from the first 45 rounds, which is statistically significant (one-sided t-test, $p < 0.001$).

In the "Self" task, which was completed by 38 subjects from the primary study, receivers switch from playing a human sender during the first 45 rounds to guessing the state from their own past non-reports. Subjects should know their own reporting strategies, so unless they have problems conditioning on actions, then subjects should be able to reliably guess the average secret number when they did not report. However, we find that receivers guess higher than their average past non-reported secret number by 0.341 on average, and this increase is statistically significant (one-sided t-test, $p = 0.002$).

The concept of conditioning failures was formalized in a behavioral approach called "Cursed Equilibrium", which was introduced by Eyster and Rabin (2005). In fact, they feature the disclosure game as a potential application of their approach. Under Cursed Equilibrium, receiver beliefs about the average non-disclosed state are a weighted average of (1) correct beliefs about the average non-disclosed state and (2) beliefs that the sender could be any type (with equal probability). Sender beliefs are correct because they have no actions to condition on.

The weights used by receivers to form beliefs are a free parameter of the model. A receiver with a weight of zero has perfectly correct beliefs about the average non-disclosed state, so will guess 1.5. A receiver with a weight of one thinks that the sender

could be of any type, so will guess 3. Thus, by varying this weight, Cursed Equilibrium can explain any receiver choices from 1.5 to 3.

A substantial majority of subjects in our experiment best respond to the distribution of receiver actions when they are in the role of sender which is consistent with Cursed Equilibrium. On the receiver side, the mode guesses of 77.4% of receivers in our primary study and 68.3% of receivers in our replication study are between 1.5 and 3, so the behavior of these subjects can be explained by Cursed Equilibrium.

The largest percentage of subjects that this theory fails to explain is the 17.5% of subjects that make a mode guess of 1 (the same percentage for both studies). However, in the “Past” task (which will be discussed later), we see evidence that some subjects may have believed the best guess was 3, but because of punishment motives, actually guessed 1.

4.4.2 Alternative Explanations: Choice Errors and Heuristics

One potential explanation for why we see overly high guesses is that receivers make choice errors with some probability. This is the reasoning behind some stochastic models of choice, such as Logit demand and Quantal Response Equilibrium.

However, a multiple factor Analysis of Variance (ANOVA) for period and subject shows that just 1.5% of the variation in guesses is due to differences in subject choices over the course of the experience, while 74.3% of the variation in guesses is due to differences in choices between subjects. Thus, it appears that overly high guess are not the result of random over-guessing, but systematic over-guessing by some subjects.

Another potential explanation for high guesses is that receivers choose the middle of the state space as choice heuristic. However, not all over-guessing is due to guesses of three, and in the “Summary Reports” task, we see that guesses of three can be changed with additional information. In addition, we might expect the response times to be shorter when subjects use heuristics (as suggested by the “Dual-Process” literature), but the response times for choices of 3 in our primary and replication studies

were not statistically different from the average response time across guesses, which was 11.0 seconds (two-sided t-test, $p=0.292$).

4.4.3 Alternative Explanation: Risk Aversion

One possible reason for variation in receiver choices is variation in risk preferences among subjects. Because the payoff function that subjects face is concave, higher risk aversion could lead receivers to guess closer to the middle of type space (a secret number of 3) than they would otherwise. Thus, such behavior could explain guesses above equilibrium or the best response.

To get a handle on the role that risk aversion plays in receiver choices, we use a standard measurement tool from experimental economics for assessing risk preferences. As described in the previous section, 38 subjects from our primary study completed the “Holt-Laury” task after the initial 45 rounds. When a subject has more than one switch point in the Holt-Laury multiple price list, then risk preference are hard to ascertain, but just 3 subjects had multiple switch points.

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

We used an OLS regression of guess onto switch point to look for evidence of positive relationship between risk aversion and the size of guesses. Controlling for the number of rounds that a receiver had spent as a sender or receiver up to that point and subject fixed effects, the coefficient on switch point is indeed positive, but is small (0.012) and not significant ($p=0.648$).

4.4.4 Alternative Explanation: Social Preferences

Another potential reason for why we observe guesses above the equilibrium prediction is that receivers may guess higher than they would otherwise to reduce the imbalance in payoffs between senders and receivers. Because of the concavity of the payoff function, when receivers make very low guesses, sender payoffs are very low. In many standard social preference models, agents lose utility when they experience guilt over making much higher payoffs than their opponent. Such models would predict that receivers would make higher guesses, even when the secret number is not reported.

For evidence of this, we examine 26 subjects from the primary study who completed the “Past” additional task. As mentioned previously, these subjects guessed the secret number from an earlier round, but without payoff implications for the sender. If social preferences were a leading explanation for higher guess, we would expect a decrease in guesses in this task. Instead, the average guess increased by 0.269, which is a statistically significant increase at a 10% level (one-sided t-test, $p=0.062$). In fact, 4 of 26 subjects switched from a guess of 1 to a guess 3. Taken together, this appears to be evidence for a punishment motive towards those who do not disclose. Instead of force pushing away from equilibrium, the social element in choice (which could be due to social preferences) appears to be pushing behavior towards equilibrium.

4.5 Robustness: Computerized Senders

One potential objection to our results is that the behavior of firms may not be well represented by a single human sender. However, this objection is mitigated by the fact that senders best respond to receiver behavior in our experiment, so how much better could a firm do?

It is possible that receivers may view firms and single human senders in different ways. We would argue that in the case of a “Mom and Pop” restaurant, the disclosure behavior of the firm is often the response of a single human sender. However, it might be that consumers view a large chain restaurant as a different kind of agent.

To provide some robustness along this dimension, we use computerized senders to approximate perceptions about large firms. In our “Computer” task, receivers switch

from playing a human sender to playing a computer sender, and as described previously, the computer plays optimally given the choices of all other receivers in the past round. This should reveal the extent to which choices are shaped by playing human senders.¹⁰

For the 34 subjects in our primary study that completed this task, we find that receiver guesses decrease slightly, but not enough to make full disclosure a best response for all computer senders. The average difference between the mode guess when playing against human senders and the mode guess when playing against the computer sender drops by just 0.103, which is not statistically significant (one-sided t-test, $p=0.128$).

Because computer senders considered the behavior of other receivers in the past round when deciding whether to report to a receiver, the likelihood of reporting was not the same for all computer senders. When the secret number was 2, computer senders disclosed the number around 56.8% of the time, close to the rate of human senders.

4.6 Informational Interventions

One way to correct for incorrect beliefs about non-disclosed types is to provide additional information to receivers and/or senders. We implement three different information interventions in the laboratory. The first to provide feedback after each round, which is the approach taken in the existing literature. The second and third are forms of aggregate information, which are described previously as the “Average Reports” task and “Consumer Reports” task.

To examine the impact of feedback after each round, we ran an additional experiment that was identical to the one described previously, but with one change: after each round, senders and receivers were told the actual secret number in that round for their pairing, the receiver’s guess, and the payoff implications of these

¹⁰ This task also changes the social considerations, but we have found these to be negligible in our experiment using the “Past” task.

choices. This additional experiment was also run in the CESS laboratory at New York University, and 34 subjects completed the experiment.

Over the last 10 rounds of this experiment, when the secret number was 3, 4, or 5, the secret number is reported 100% of the time. Also, in contrast to our results without such feedback, a secret number of 2 is reported 85.2% of the time. A secret number of 1 is still unlikely to be reported, as it is reported 28% of the time, but this does not go against the equilibrium predictions.

In the last 10 rounds, receiver guesses were also close to the equilibrium predictions. In rounds where the secret number is not reported, 81.8% of guesses are 1, with an average guess of just 1.227.

In the “Average Reports” and “Consumer Reports” tasks, it is possible to observe both the direct impact of the additional information and the strategic impact of the additional information. After 45 rounds, the subjects are shown the corresponding aggregate information, and then guess the non-reported secret number from a previous round, as in the “Past” task. Because the strategy of the sender is held fixed, we can see the direct impact of the information. We then have subjects play 5 more rounds as before, but now the sender’s strategy can adjust – the sender knows the information and knows the receiver knows the information too – so this part identifies the strategic impact of the additional information.

The effect of information in the “Average Reports” task on X subjects from the primary study, while slight, is actually to move receiver guesses further away from the best response. Looking first at the direct effect, the guess of a past non-reported secret number increases over the mode guess from past rounds by 0.117 on average, but this increase is not statistically significant (one-sided t-test, $p=0.115$). However, the addition of information did move guesses further from the best response. Before the information, the mode guess from past rounds is 0.55 away from the best response on average, and after the information, the guess of a past round is 0.7 away from the best response on average. This increase is statistically significant at a 5% level (one-sided t-test, $p=0.0132$).

Despite the deleterious direct impact on this information on beliefs, the rate of disclosure for secret numbers of 2 actually increases slightly. In the 5 additional rounds, the reporting rate of a secret number of 2 increases by 6.1 percentage points, but this increase is statistically insignificant (one-sided t-test, $p=0.3316$).

On the other hand, the effects of information in the “Consumer Reports” task on 34 subjects from the primary study are to move choices closer to the best response and to substantially increase disclosure – leading choices much closer to equilibrium. Looking first at the direct effect, the guess of a past non-reported secret number decreases from the mode guess from past rounds by 0.171 on average, and this decrease is statistically significant (one-sided t-test, $p=0.0482$).

In addition, the distribution of guesses shifts towards the best response. Before the information, the mode guess from past rounds is 0.724 away from the best response on average, and the after the information, the guess of a past round is 0.553 away from the best response on average. This decrease is statistically significant at a 5% level (one-sided t-test, $p= 0.0108$). The outcome of these changes is shown in Figure 8, which shows the histogram of the mode of guesses before the intervention and guesses of a past round after the informational intervention.

Finally, in the 5 additional rounds, the reporting rate of a secret number of 2 increased to 73.9% from a reporting rate of 45.5% in the previous 40 rounds. This increase of 28.5 percentage points is statistically significant (one-sided t-test, $p=0.005$). However, this rate did not increase appreciably over the 5 additional rounds, suggesting that this information intervention is not enough to produce convergence to full disclosure. From the first two rounds to the next three rounds, the reporting rate of a secret number of 2 increased just 6.9 percentage points, which was not a statistically significant increase ($p=0.3615$).

5 Robustness: Enlarged State Space

In many settings of verifiable disclosure, the size of the state space is under the control of policy makers. For instance, restaurant hygiene can be reported on a scale of

1-5 (as it is in some areas of the UK) or even 1-100 (as it is in some areas of the US). To get a sense for how the size of the state space might impact our findings, we ran an experiment that was similar to our primary experiment, but with an enlarged state space.

In this new experiment, the state space is $S=\{1,2,3,4,5,6,7,8,9,10\}$, which is twice as large as in the original experiment. Here again we allow receivers to guess half-unit intervals, so the action space is $A=\{1,1.5,2,2.5,\dots,9,9.5,10\}$.

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

Aside from increasing the set of secret numbers and changing the payoff table, the experimental design and instructions are the same as in the original experiment. We conducted this experiment in the same location as our replication study, which is the Computer Lab for Experimental Research (CLER) facility at the Harvard Business School. 84 subjects completed the new experiment, and the median and mode session size was 14 subjects. In order to make this new experiment more directly comparable to our primary study, we limited the subject pool to be 25 years old or younger.

5.1 Results

Figure 9 shows the average reporting rate by secret number in the new experiment. As in the primary study with 5 secret numbers, the reporting rate increases monotonically with the secret number. The reporting rate for a secret number of 3 in the new experiment is 41.8%, which is comparable to the reporting rate for a secret number of 2 in the primary study of 40.8%. They are also comparable in the sense that a risk neutral sender is close to indifferent between reporting and not reporting at a secret number of 3 in the new experiment. The average guess for a non-reported secret

number is 3.239 with a 95% confidence interval of 3.063 to 3.416. In addition, the reporting rate in the new experiment for a secret number of 5 is 85.7%, which is similar to the reporting rate 88.6% for a secret number of 3 in the primary study.

As Figure 10 shows, there is also heterogeneity in receiver responses in the new experiment: 12.1% of subjects make the equilibrium guess of 1, 19.0% of subjects (the largest percentage) guess 2, 13.5% guess 2.5, and 11.3% guess 5. For a risk neutral receiver, the best response is to guess the average non-reported secret number, which is 2.503 with a 95% confidence interval of 2.353 to 2.653. This is far below the average guess from receivers in the experiment (3.239 with a 95% confidence interval of 3.063 to 3.416). Once again, we find evidence of little learning or stochasticity in choice for receivers. Using an ANOVA test, we find that just 3.7% of the variation in guesses occurs over rounds, while 72.8% of the variation in guesses occurs across subjects.

5.2 Informational Interventions

In half of the sessions of the new experiment, subjects completed the “Consumer Reports” additional task, and in the other half subjects completed the “Average Reports” additional task. Once again, we found the former able to move beliefs, but not the later.

Looking first at the direct effect of “Average Reports”, we find that the difference between the mode guess and the best response is 1.211 before the information, and 1.184 after the intervention, which is not significantly different (two-sided t-test, $p=0.895$). However, the “Consumer Reports” intervention is able to improve guesses substantially. The difference from the best response is 1.544 before the intervention and 1.111 after the invention, which is significantly different (two-sided t-test, $p=0.001$)

6 Discussion and Conclusions

Our findings demonstrate that - in contrast with the unraveling hypothesis - full disclosure does not occur even in an extremely simple situation in which all of the institutional assumptions of unraveling are satisfied. The failure stems from the fact that

receivers do not infer that no news is bad news, despite knowing the distribution of quality scores and the fact that senders have chosen not to reveal information.

Our findings cannot be explained by the typical disclosure theories with rational agents. By design, we removed possibilities such as disclosure costs, strategic concerns, sender claims not to know the true state, and receiver's lack of information on potential states from our experimental design. We also rule out random choice, risk aversion, and social preferences as potential explanations.

Our findings generate potential remedies to the failure of unraveling. As shown in Forsythe, Isaac and Palfrey (1989), Forsythe, Lundholm and Rietz (1999), and in our additional experiments, full disclosure occurs when the experimenter shows the receiver the true state after each round of her guess. However, such a full feedback remedy is very difficult to implement in reality, as it takes time and effort for a third party to find out the truth and it is costly to convey the truth to the receiver in a precise and timely manner.

A more practical remedy is to provide a summary of reported states (like our "Average Reports" task) or a summary of reported and unreported states (like our "Consumer Reports" task). As shown above, the "Average Reports" intervention has little direct effect on beliefs. In fact, it makes some receivers even more optimistic than before about non-disclosing senders. Also, our results suggest that an "Average Reports" intervention fails to move the market towards unraveling. Ironically, more information by the "Average Reports" intervention may have the potential to reduce the transparency of voluntary disclosure.

In contrast, the "Consumer Reports" intervention makes receivers more skeptical about non-disclosure and therefore pushes the market towards the unraveling equilibrium. This confirms our conclusion that the failure of unraveling is driven by receivers' incorrect beliefs about non-disclosure. It also suggests that voluntary disclosure must be accompanied by extra information about non-disclosed states in order to foster full disclosure. However, this intervention did not lead to full disclosure in the laboratory. The optimal intervention will depend on weighing the cost of

producing extra information and the benefits of closing the gap in beliefs and disclosure rates.

Our study has several limitations. The subjects in our main experiments are undergraduate students at New York University, who may not be representative of the average population targeted by a voluntary disclosure policy. However, we might expect people with some college education to be more sophisticated than the average population in the US. This suggests that the failure of unraveling shown in our data may be a conservative estimation of what would happen in a real market. In addition, we are able to replicate our results with a subject pool that contains a mixture of ages at another university, suggesting that the failure to perceive no news as bad news is widespread.

In addition, human subjects may not act as a firm does, which may cause some to question the generalizability of our results. However, in many settings disclosure decisions are undertaken by small firms, such as sole proprietorships, so human subjects may be suitable replacement for firms. For example, one might think of a family-run restaurant deciding whether to disclose its hygiene inspection grade. In addition, the results from our “Computer” task indicate that failures of unraveling can occur even when senders are not susceptible to any behavioral biases or mistakes. Although there is evidence that large firms do not perfectly maximize profits, it seems likely that their choices lay somewhere between the decisions of a single undergraduate student and a perfectly optimizing computer.

The third limitation is that, by design, we do not know how bad inferences interact with other rational factors that impede full disclosure. Will bad inferences have a greater impact on disclosure rate when disclosure is costly or when senders have a choice to learn the true state? Will the “Average Reports” and “Consumer Reports” interventions motivate more disclosure when disclosure is costly? These questions warrant future research.

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. B., Jin, G. Z., Leslie, P., Quinn, A. J., & Zou, B. (2014). Incomplete Disclosure: Insights from Maricopa. Mimeo.
- Benndorf, V., Kübler, D., & Normann, H. T. (2015). Privacy concerns, voluntary disclosure of information, and unraveling: An experiment. *European Economic Review*, forthcoming.
- Bhattacharya, S., Kang, J. W., & Wilson, A. (2015). Cheap Talk and Persuasion: Theory and Behavior. Mimeo.
- 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.
- Cai, H., & Wang, J. T. Y. (2006). Overcommunication in strategic information transmission games. *Games and Economic Behavior*, 56(1), 7-36.
- Dickhaut, J., Ledyard, M., Mukherji, A., & Sapro, H. (2003). Information management and valuation: an experimental investigation. *Games and Economic Behavior*, 44(1), 26-53.
- DellaVigna, S., & Gentzkow, M. (2009). Persuasion: empirical evidence (No. w15298). National Bureau of Economic Research.
- Dranove, D., & Jin, G. Z. (2010). Quality Disclosure and Certification: Theory and Evidence. *Journal of Economic Literature*, 48(4), 935-963.
- 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.
- 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: 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.
- Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, 92(5), 1644-1655.
- Ivanov, A., Levin, D., & Niederle, M. (2010). Can relaxation of beliefs rationalize the winner's curse?: an experimental study. *Econometrica*, 78(4), 1435-1452.
- Jin, G. Z. (2005). Competition and disclosure incentives: an empirical study of HMOs. *RAND Journal of Economics*, 93-112.
- Kessler, J. B., & Roth, A. E. (2012). Organ Allocation Policy and the Decision to Donate. *American Economic Review*, 102(5), 2018-47.
- King, R. R., & Wallin, D. E. (1991). Voluntary disclosures when seller's level of information is unknown. *Journal of Accounting Research*, 96-108.
- Luca, M., & Smith, J. (2013). Salience in quality disclosure: evidence from the US News college rankings. *Journal of Economics & Management Strategy*, 22(1), 58-77.

- 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. *The RAND Journal of Economics*, 18-32.
- Niederle, M., & Vesterlund, L. (2007). Do women shy away from competition? Do men compete too much? *Quarterly Journal of Economics*, 122(3), 1067-1101.
- Sánchez-Pagés, S., & Vorsatz, M. (2009). Enjoy the silence: an experiment on truth-telling. *Experimental Economics*, 12(2).
- Serra-Garcia, M., van Damme, E., & Potters, J. (2011). Hiding an inconvenient truth: Lies and vagueness. *Games and Economic Behavior*, 73(1), 244-261.
- 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: Fraction of rounds each secret number was reported (with 95% confidence intervals)




Figure 2: Fraction of rounds each secret number was reported in the replication study (with 95% confidence intervals)




Figure 3: Histogram of receiver guesses of non-reported secret numbers




Figure 4: Histogram of receiver guesses of non-reports secret numbers in the replication study




Figure 5: The average guess of non-reported secret numbers by block of 5 rounds (with 95% confidence intervals and the average guess across all rounds)

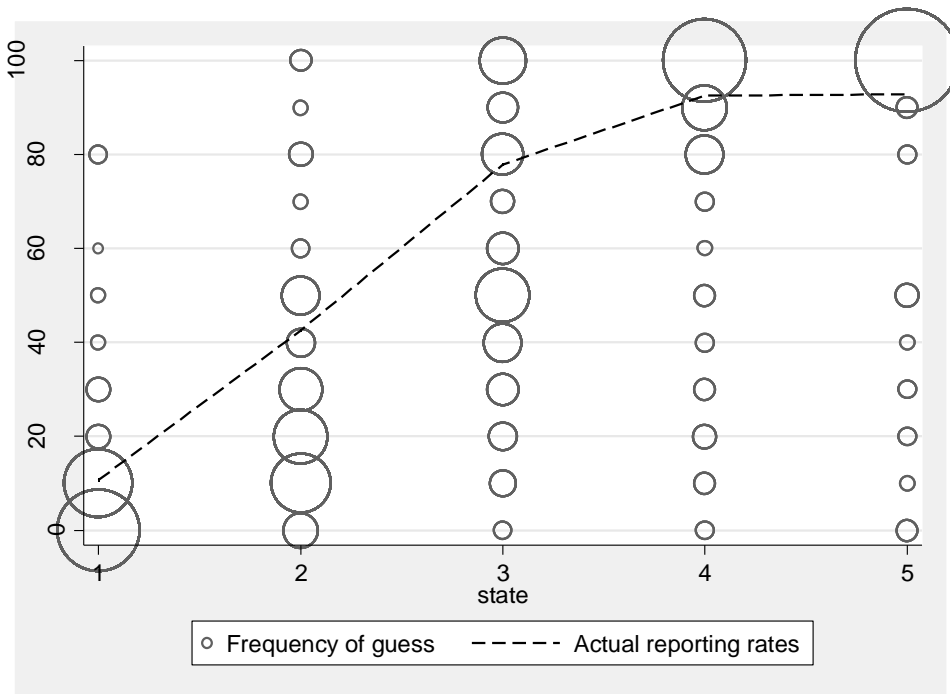


Figure 6: The frequency of guesses for the reporting rate at each secret number (with the actual reporting rate at each secret number)

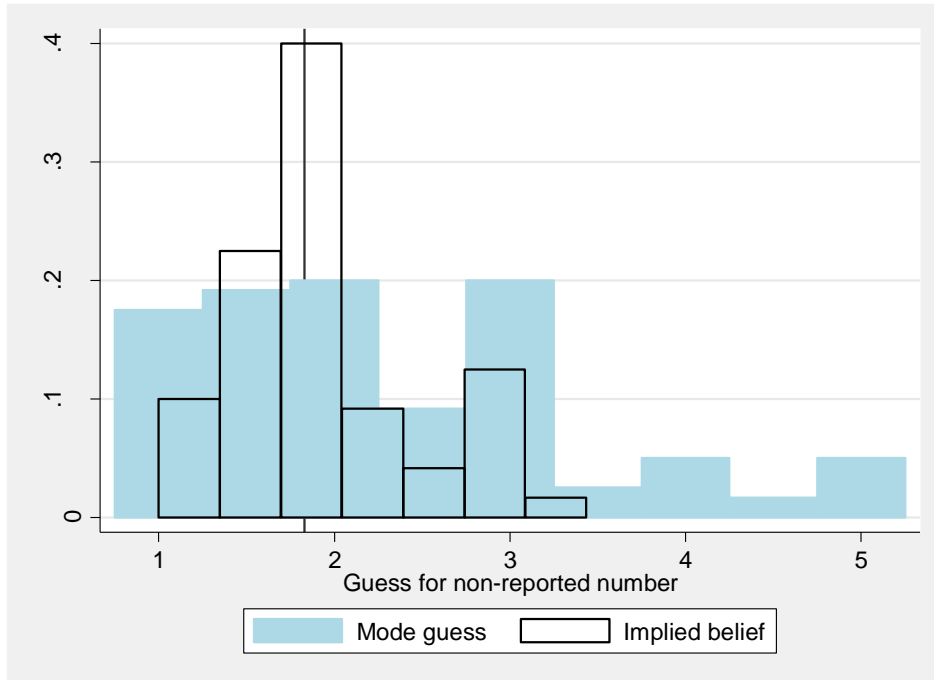


Figure 7: The implied belief of the average non-disclosed number and the average guess of the non-disclosed number (with the average non-disclosed number)

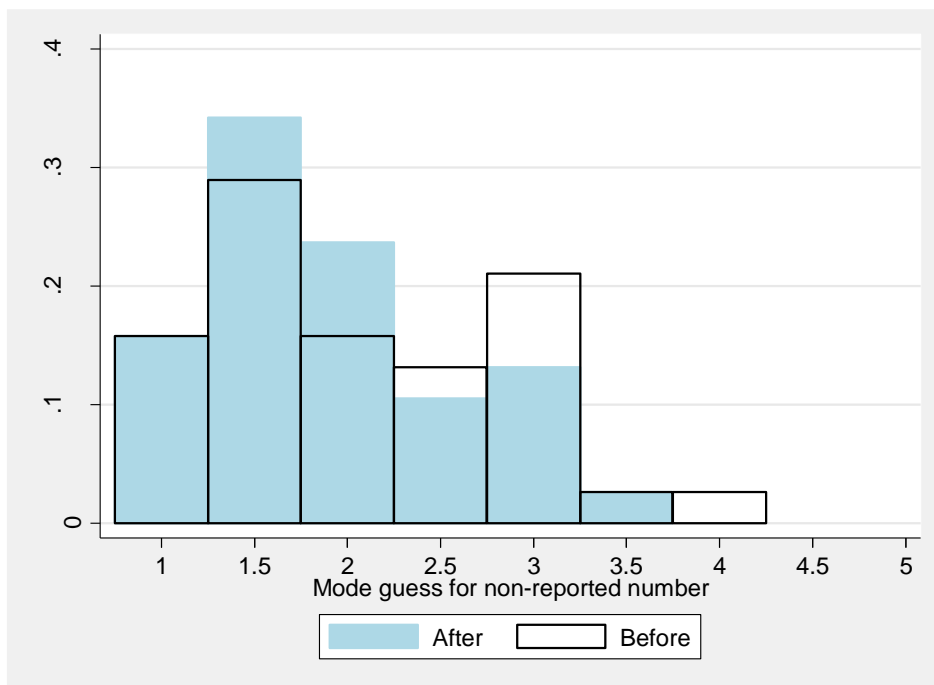


Figure 8: The histogram of mode guesses of non-reported secret before the “Consumer Reports” intervention and of the guesses of a past round after the intervention




Figure 9: Fraction of rounds each secret number was reported in the robustness experiment (with 95% confidence intervals)




Figure 10: Histogram of receiver guesses of non-reported secret numbers in the robustness experiment

Appendix: Experimental Instructions

Welcome

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

Please silence and put away your cellular phones now.

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

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

Instructions

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

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

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

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

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

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

Appendix: Experimental Instructions

PAYOFFS: S , R	R's guess: 1	R's guess: 2	R's guess: 3	R's guess: 4	R's guess: 5
Secret number: 1	-29 , 110	17 , 90	57 , 57	90 , 17	110 , -29
Secret number: 2	-29 , 90	17 , 110	57 , 90	90 , 57	110 , 17
Secret number: 3	-29 , 57	17 , 90	57 , 110	90 , 90	110 , 57
Secret number: 4	-29 , 17	17 , 57	57 , 90	90 , 110	110 , 90
Secret number: 5	-29 , -29	17 , 17	57 , 57	90 , 90	110 , 110