

Complex Disclosure

Ginger Zhe Jin¹

University of Maryland and NBER

Michael Luca

Harvard Business School

Daniel Martin

Northwestern Kellogg School of Management

February 18, 2021

Abstract

We present evidence that unnecessarily complex disclosure can result from strategic incentives to shroud information. In our lab experiment, senders are required to report their private information truthfully, but can choose how complex to make their reports. We find that senders use complex disclosure over half the time. This obfuscation is profitable because receivers make systematic mistakes in assessing complex reports. Regression and structural analysis suggest that these mistakes could be driven by receivers who are naive about the strategic use of complexity or overconfident about their ability to process complex information.

¹ Part of the research was conducted when Jin took leave at the Federal Trade Commission. The views expressed are those of the authors and do not necessarily represent those of the U.S. Federal Trade Commission or any individual Commissioner. Martin would like to thank both the Paris School of Economics and the Camargo Foundation for their hospitality during the writing of this paper. Early stages of this project were supported by the French National Research Agency, through the program Investissements d'Avenir, ANR-10--LABX_93-01. We would like to thank Patrick Rooney, Byron Perpetua, and Philip Marx for excellent assistance. All rights reserved. All errors are ours.

1. Introduction

Firms are often required to disclose contract terms and other relevant information to consumers. For example, credit card companies are required to disclose interest rates. Tech companies are required to disclose privacy policies. And public firms are required to disclose financial performance. Ben-Shahar and Schneider (2014) and Lowenstein et al. (2014) provide a long list of hard-to-understand disclosures and have argued that mandated disclosure has failed as public policy as a consequence. Financial experts even blame the complexity of financial products for the 2008 financial crisis, although the risks embedded in these products were supposedly disclosed to a ratings agency.²

Some of these disclosures are complex by necessity – simply because firms need to provide very detailed information. However, because firms have the ability to manipulate the complexity of their reports, some of these disclosures may be far more complex than they need to be. For example, credit card companies can present payment schedules, penalties, and fees clearly or bury potentially important details in the fine print.³ Privacy policies can be written in easy to understand language or shrouded in pages of complex legalese. When public firms make financial disclosures, they can summarize them in several paragraphs or run them as long as 257 pages.⁴

The extent to which firms can exploit consumers by increasing complexity depends crucially on how consumers respond when they observe complex information. If consumers are sufficiently skeptical about firms that use complex disclosures and account for this in their decision-making process, then firms that offer better terms or higher-quality products will want to present this information as clearly and simply as possible to prevent themselves from being mistaken for worse firms. As a result, we would expect only the worst firms to use complex disclosures, which is similar to the “unraveling” results in voluntary disclosure (Viscusi 1978;

² <https://www.ft.com/content/24f73610-c91e-11dc-9807-000077b07658> accessed on September 26, 2017.

³ The Truth in Lending Act of 1968 (TILA) requires lenders to disclose consumer credit terms and cost in a standardized way. The Real Estate Settlement Procedures Act of 1974 (RESPA) requires lenders and others involved in mortgage lending to provide borrowers with pertinent and timely disclosures regarding the nature and costs of a real estate settlement process. In 2015, the US Consumer Finance Protection Bureau consolidated the disclosure requirements under TILA and RESPA, resulting in the Loan Estimate Form and the Closing Statement Form, which standardize the content and format of disclosure in mortgage lending.

⁴ The SEC does not impose a limit on the length of a filing, and the average 10-K has grown from roughly 30,000 words in 2000 to 42,000 words in 2013, with GE’s 2014 10-K stretching to 103,484 words and 257 pages. Source: <https://www.wsj.com/articles/the-109-894-word-annual-report-1433203762>.

Grossman and Hart 1980; Grossman 1981; Milgrom 1981). However, systematic mistakes by consumers when facing complex reports can give rise to strategic incentives to complexify, motivating companies to choose complexity over simplicity in their disclosures.

In reality, it is difficult to determine if consumers are sufficiently skeptical about firms that use complex disclosures and account for this in their decision-making process, because it is hard to discern whether disclosures are complex by necessity or the outcome of firms strategically choosing to make information unnecessarily complex. To overcome this difficulty, we design a laboratory experiment to study the strategic use of complexity in a controlled setting. In our experiment, complexity arises only from its strategic use, the conflicting interests of senders and receivers are clear, the amount and over-use of complexity are quantifiable, and the beliefs of agents are easily elicited.

There are two roles in our experiment: a sender (e.g., the firm) and a receiver (e.g., the consumer). Subjects are randomly paired in each round, one randomly assigned to be the sender and the other to be the receiver.⁵ In each round, the sender observes a new state (which is an integer drawn uniformly from 1 to 10) and chooses how complex to make their report of the state. Based on the sender's report, the receiver then makes a guess of the state.

We impose a strong and clear conflict of interest: the sender would like receivers to guess that the true state is as high as possible, and the receiver would like to guess as accurately as possible. In our main sessions, we debrief both players at the end of each round about the true state, the sender's choice of complexity, and the receiver's guess in that round. This way, subjects have many opportunities to learn the strategic forces in the game and the consequences of their actions.⁶

When a sender's report is simple, the state is presented as a single integer. When their report is complex, the state is presented as several computer-generated numbers (up to 20) that add up to the state. While this is just one of many ways to operationalize complexity, it has the advantage that individuals have experience with the task and may hold well-formed beliefs about their ability to internalize complexity of this form. In

⁵ Roles were randomly assigned so that subjects could experience both roles, which allowed subjects to be well informed about the actions and payoffs available in both roles.

⁶ As a robustness check, we also ran sessions without feedback. In addition, we ran a robustness check where we limit the number of complexity levels available to senders. See Section A.4 of the appendix.

addition, this task allows us to easily measure the extent of complexity and the size of the sender and receiver mistakes.

Although senders can make their reports complex, they cannot misrepresent the underlying state, as the computer-generated numbers will always add up to the state. Thus, the theoretical predictions for behavior in our experiment are stark. If complexity generates sufficient uncertainty about the state, it is easy to show using the unraveling arguments mentioned previously that the unique sequential equilibrium is full use of simple reports for all states above the worst, even though there is a conflict of interest between senders and receivers in our experimental game.

Instead, defining “low complexity” as messages with 5 or fewer numbers to sum, senders use low complexity less than 50% of the time, even if we look just at the second half of rounds. When not using low complexity, senders mostly use high complexity (defined as messages with 15 or more numbers to sum), and they do so in a systematic way.⁷ Over rounds, senders gravitate towards two extremes: using low complexity for high states (secret numbers of 8 and up) and using high complexity for low states (secret numbers of 3 or less). When the state is neither high nor low, senders use high complexity approximately 33% of the time, even in the second half of rounds.

Why is complex disclosure so prevalent? One possibility is that senders use complex disclosure more than they should. However, we find that using high complexity to hide both low and middle states is profitable for senders because receivers guess higher than the actual secret number in both low and middle states, and this persists into the second half of rounds.⁸ Because sender behavior is largely consistent with these strategic incentives, the average losses of senders are small and decrease over rounds.⁹

So why do receivers systematically guess higher than the actual secret number when disclosures are complex? The driving force of complexity in disclosure is that it adds noise to messages because receivers are unable to fully internalize complex information. Because of this noise, receivers face an involved decision-making process:

⁷ Senders use middle complexity at a similar rate across states, so the only variation across states is the frequency that low and high complexity are used.

⁸ We define the “optimal” action for senders as the one that has the highest expected payoff, and we measure losses in terms of expected payoff.

⁹ The primary sources of losses are choosing high complexity at high states and choosing low complexity at low states. These mistakes, along with their possible sources, are examined in Section 4.2.

they must infer why complexity was used and how much noise complexity added to the message, and then adjust their guesses accordingly.

There are several ways that this process can be distorted. For example, a receiver who is naive about the strategic use of complexity will not adequately adjust their guesses for the bad news communicated by complexity, which will lead them to systematically guess above the actual secret number. In fact, naivete is the leading assumption that justifies the strategic use of complexity in theoretical models (for example, see Gabaix and Laibson 2006; Spiegel 2006; Carlin 2009; and Armstrong and Vickers 2012), and it has been found to impact other forms of disclosure (Cai and Wang 2006; Jin, Luca, and Martin forthcoming).

However, naivete is not the only reason why receivers might not adequately adjust their guesses for the bad news communicated by complexity. A receiver who is overconfident about their ability to understand complex disclosures will overly trust their reading of a complex report when figuring out what action to take. In our experiment, overconfidence about the ability to distill complex information can lead receivers to ignore the possibility their reading of a complex report is wrong and that the secret number is in fact much lower, which is likely to be the case when reports are complex. For example, if a receiver quickly adds up the numbers in a complex report to 7, they should be wary about guessing 7 because secret numbers that high are very unlikely when the message is complex. However, if they are overconfident about their math ability, they will be less worried about how unlikely the secret number is to be 7 when the report is complex. Thus, overconfidence in their ability to digest complex information can lead receivers to underweight the fact that the information being complex is bad news – even when they hold correct strategic beliefs about complex messages.

To study the impact of naivete and overconfidence on receiver guesses, we elicit beliefs from subjects about the strategic implications of complexity and about their math ability. 12.6% of subjects appear naive about complexity because they guess that the average secret number was higher than it actually was when reports were complex. 33.8% appear overconfident about their math ability because they guess that they performed better on a math test than they actually did. Moreover, when receivers are

either more naive about complexity or more overconfident about their math ability, they systematically over-guess by a larger amount.

We expand on these results with a structural analysis. Because the impact of naivete and overconfidence on receiver guesses is indirect and distorted by boundary effects, it is hard to disentangle the extent to which receiver mistakes can be attributed to naivete and overconfidence. To investigate, we use a partial-equilibrium structural model in which receivers observe a level of complexity, apply their strategic beliefs about the link between complexity and states, receive a noisy signal of the state based on the complexity level (due to math errors), update their beliefs about the state, and then make a guess. Importantly, we estimate math errors *out-of-sample* using a math test in which subjects face high complexity without any strategic considerations.

As a baseline, we close the model by assuming that strategic beliefs are correct and belief updating is perfect, and this version of the model fails to predict the over-guessing we observe at middle states and the extent of under-guessing we observe at high states. If we assume instead that 12.6% of receivers are “Level-1” naive (believe that each state is equally likely when reports are complex), we are better able to predict receiver guesses, but this assumption does not generate enough naivete to produce over-guessing of middle secret numbers. However, adding overconfidence about math ability at the levels we observe in a math test generates even more accurate predictions of receiver guesses and does produce over-guessing of middle secret numbers. This structural analysis demonstrates that a simple behavioral model with naivete and overconfidence can largely explain the choices we observe in our experiment.

Evidence that strategic disclosure can be impacted by naivete is consistent with results from analogous voluntary disclosure experiments in which senders can either disclose or not disclose their information. However, we find two main differences between non-disclosure and complex disclosure. First, complexity persists and is effective even with repeated feedback, unlike non-disclosure. Second, we find evidence that complex disclosure is impacted by an additional bias: overconfidence about ability.

This factor is unique to complex disclosure, where noise in messages arises from limits in ability, not external randomization devices.¹⁰

To our knowledge, this is the first lab experiment to examine the strategic use of complexity for a simple sender-receiver framework in which preferences are known. It provides evidence from a controlled environment that complex disclosure can result from strategic incentives to shroud information. In addition, our results suggest that strategic naivete and overconfidence about the ability to internalize complexity might be important drivers of these incentives. Naivete has been proposed in the theoretical literature as reason for the strategic use of complexity, but to the best of our knowledge, overconfidence about ability has not been previously proposed as an explanation for complex disclosure.¹¹ Also, our structural estimates based on naivete and overconfidence add to a growing literature that imposes model structure to help understand the relative importance of different behavioral biases (DellaVigna 2018).

The mechanisms we identify in our experiment may most directly help to shed light on such as insurance, credit card, and investment choices, where disclosure frequently is complex and involves math calculations. This relates to a growing body of field evidence demonstrating that many consumers make systematic mistakes when making such calculations. For example, Bhargava, Loewenstein and Sydnor (2017) examine the health plan choices of 23,894 employees at a U.S. firm, based on a large menu of options that differed only in financial cost-sharing and premium. They find that the majority of employees chose dominated plans, which resulted in excess spending that is equivalent to 24% of chosen plan premiums. Similarly, consumers that responded to a lender's *inferior* solicitation of preapproved credit card offers (Agarwal et al. 2010), and credit card issuers provide offers with back-loaded and hidden features, upfront rewards, visual distractions, and fine print at the end of the offer letter (Ru and Schoar 2016). Systematic mistakes are also found in consumer choice of mortgage loans (Agarwal et al. 2017) and pension funds (Duarte and Hastings 2012). While the calculations made in

¹⁰ Overconfidence about ability is less relevant when noise is added mechanically, such as when messages are randomly selecting from an interval (as in Cai and Wang 2006).

¹¹ Grubb (2015) presents evidence of how other forms of overconfidence interact with complex disclosures, such as overconfidence about the precision of estimates, overconfidence about self-control, and overconfidence about attention to fulfilling contract terms.

these settings are far more challenging than the addition task faced by our subjects, this might make overconfidence more of a factor, as absolute overconfidence has been found to increase with task difficulty (see Moore and Healy 2008). More broadly, our results speak to a phenomenon that occurs in a variety of disclosure contexts. People are routinely given complex information to digest. Our results suggest there may be a strategic component to the level of complexity, as well as systematic mistakes in the inferences made based on those complex disclosures.

The rest of the paper is organized as follows. Section 2 reviews three related literatures and articulates our contribution to each. Section 3 presents our experimental design. Section 4 discusses our experimental results, looking both at overall behavior and dynamic patterns using summary statistics, regression analysis, and structural estimation. Section 5 concludes with potential policy implications.

2. Literature Review

Our paper draws on and contributes to three literatures: the literature on voluntary and mandatory disclosure, the literature on obfuscation and behavioral biases, and the experimental literature.

2.1 Voluntary and Mandatory Disclosure

In virtually every transaction imaginable, companies must decide what information to disclose. In practice, voluntary disclosure is observed in many industries, but is far from complete.¹² As summarized in Dranove and Jin (2010), this incompleteness can be explained by external factors such as disclosure cost and consumer knowledge before disclosure or by a seller's strategic incentives.¹³

¹² See Mathios (2000), Jin (2005), Bollinger et al. (2011), Bederson et al. (2018), Anderson et al. (2015), Fung et al. (2007), and Luca and Smith (2015) for specific examples.

¹³ For instance, see Jovanovic (1981) for the impact of disclosure cost on disclosure decisions, see Matthews and Postlewaite (1985) on the incentive to not knowing true quality, see Fishman and Hagerty (2003) for the impact of having some buyers not understand disclosures, see Feltovich et al. (2002) on relating disclosure to counter-signaling, see Board (2009) on the incentive to use disclosure for differentiation, see Grubb (2011) on the incentive to hide due to dynamic concerns, and see Marinovic and Varas (2016) on disclosure decisions in light of litigation risk.

Even if the disclosure itself is mandatory, firms can often choose the content or format of their disclosures. This leads to a mix of voluntary and mandated elements in reporting. For instance, policies that mandate disclosure on a limited number of dimensions may encourage firms to redirect resources to the mandated dimensions but shirk on other dimensions (Lu 2012). Even on the mandated dimensions, firms may game the definition of the mandated statistics (Dranove et al. 2003; Jacob and Levitt 2003) or shroud it in a way that obfuscates important details (Brown et al. 2010).

Because we focus tightly on the choice of simple or complex disclosure, we exclude many other external factors that could complicate a firm's choice of simplicity in a mandatory disclosure setting, such as sender uncertainty, legal concerns, and disclosure costs. In doing so, we simplify the strategic interaction between sender and receiver, which helps us to isolate how actions are driven by information and beliefs.

2.2 Obfuscation and Behavioral Biases

The empirical literature has documented several examples of obfuscation. Brown, et al. (2010) show that shipping and handling cost is often shrouded on e-commerce platforms. Sullivan (2017) shows that some hotels keep mandatory resort fees separate from room rate, and some online travel platforms conduct the price search by room rate only and do not disclose resort fees until consumers reach the hotel-specific page before payment. Obfuscation can also appear in a more sophisticated way. Ellison and Ellison (2009) document a loss-leader strategy by Internet retailers. In that strategy, the retailer sets a low price for a low-quality product on a price comparison site, and then persuades consumers to buy higher-quality products at a greater markup after consumers visit the retailer's website. Célérier and Vallée (2017) find that banks offer retail investment products in ways that are consistent with strategic obfuscation. For instance, more complex products are more expensive and are more harmful for consumers.

While these studies largely focus on the seller's choice of obfuscation, other empirical studies document the behavior of information receivers. Chetty et al. (2009) study two price regimes that include or exclude tax in the list price (tax rate is well known). They find that people are much less responsive to tax in the second regime

because taxes are more complicated to compute. Blake et al. (2017) study an online ticket platform that switched from transparent pricing to hiding transaction fees until payment. They find that consumers are more likely to buy more tickets and pay higher price if transaction fees are “back-end.” Pope (2009) and Luca and Smith (2013) show that the salience of quality disclosure determines the extent to which customers respond. In a variety of settings, people are found inattentive to relevant details even after disclosure occurs (Armstrong and Chen 2010; DellaVigna and Pollet 2005; DellaVigna and Pollet 2009; Lacetera et al. 2012; Englmaier et al. 2017). In a similar spirit, Hanna, Mullainathan, and Schwartzstein (2014) show that often consumers only attend to certain once-overlooked information when information is presented in a summary form.

Our lab experiment complements this field work by jointly studying the decisions of senders and receivers in an environment where we control both incentives and information and remove non-behavioral reasons for complex disclosure. By measuring sender and receiver mistakes at the same time, we can accurately determine departures from equilibrium and shed light on the extent to which our subjects behave optimally in response to their opponents’ actions.

In addition, because we measure subjects’ beliefs and study the extent and nature of their belief biases, our work speaks to a growing literature that models the relationship between firm obfuscation and consumer naivete. For example, Ellison (2005) shows that add-on pricing can be rationalized if one adds a subpopulation of irrational consumers. Gabaix and Laibson (2006) develop a model in which firms can shroud dimensions of product information when some consumers are myopic or unaware. Heidhues et al. (2016) further give out the conditions under which a shrouding equilibrium arises when naive consumers ignore add-on prices until at least one firm unshrouds (reveals) the additional price. Spiegler (2006) assumes consumers are only capable of evaluating one of many dimensions of the product, which motivates firms to obfuscate by making the product more attractive on some dimensions but less attractive on others. Similarly, Armstrong and Vickers (2012) model bank overdraft fees in a market where some consumers are sophisticated, and some consumers are naive, and they show that competition may end up subsidizing the sophisticated at the expense of the naive. Bianchi and Jehiel (2015) capture complexity choice in financial disclosures by allowing firms to

add noise in the signals that disclosure provides, and in their model, investors make mistakes with noisy signals because they over-extrapolate from the limited number of signals that they receive. Carlin (2009) presents a model in which firms use complex pricing when enough consumers are myopic. Hirshleifer and Teoh (2003) consider the impact of naivete on financial disclosures, where receivers can be naive about non-disclosed information and inattentive to disclosed information. These results build on a strong tradition in modeling naivete in behavioral economics, including the theories of Cursed Equilibrium (Eyster and Rabin 2005), Analogy-Based Expectation Equilibrium (Jehiel 2005), Level-k reasoning (Crawford and Iriberry 2007), and coarse thinking (Mullainathan et al. 2008).

Theoretically, receiver naivete is not a necessary condition for senders to choose obfuscation. Firms may still engage in obfuscation even if all information receivers are rational. In a model where consumers must spend time to search for price, Ellison and Wolitzky (2012) show that firms have incentive to increase consumer's search cost through obfuscation. In doing so, obfuscation increases the search cost of consumers, raises equilibrium price, and benefits all firms even if some firms do not use obfuscation themselves. In a different setting, Perez-Richet and Prady (2012) consider obfuscation to a third-party certifier (say bond rating agencies), whose job is to digest and certify the disclosed information. They find that even good types may add complexity to disclosed information because this could motivate the certifier to lower its validation threshold. de Clippel and Rozen (2020) consider a sender whose desired strategy is only implemented with an exogenous probability (the "precision level of communication") and a receiver who must exert costly effort to learn about the state when messages are obfuscated. This combination of forces generates enough uncertainty about obfuscated messages to completely stop unraveling in undominated equilibria. The testable conditions for this equilibrium are the testable conditions for a receiver who is rationally inattentive given the strategically correct prior beliefs about the state when messages are obfuscated.

2.3 Lab Experiments

Our experimental design is related to the cheap talk experiments of Cai and Wang (2006) and the voluntary disclosure experiments of Jin et al. (2018).¹⁴ For instance, we also frame states as “secret numbers” and use a similar payoff structure. The key difference in our experimental design is that senders must truthfully reveal their type and can make their reports complex. Hence, our experiment examines complex disclosure, rather than cheap talk or voluntary disclosure. This difference is meaningful: cheap talk and complex disclosure have opposite theoretical predictions when there is a strong conflict of interest between senders and receivers, and complexity introduces a new type of internally-driven noise about the state (based on limited ability) that is not present in cheap talk or voluntary disclosure.

While our experiment focuses on mandatory disclosure rather than voluntary disclosure, the choice of simplicity is voluntary and subject to the same unraveling logic. Though unraveling has been confirmed by multiple disclosure experiments, Jin et al. (2018) show that immediate and repeated feedback is crucial for subjects to converge to the predictions of unraveling. Our results suggest that in a setting different from the classical game of voluntary disclosure, even immediate and repeated feedback is not enough to salvage the unraveling prediction.

This paper also joins a growing number of experiments that study the impact of complexity on strategic interactions. For example, Sitzia and Zizzo (2011) implement a lab experiment in which sellers set the price of products of exogenous complexity, where complex products are compound lotteries. Sellers in their experiment do not offer higher prices when complexity is higher, but buyers appear to buy a higher quantity of complex products. Kalaycı and Potters (2011) implement a lab experiment where sellers determine both the price and complexity of products of exogenous quality. In their experiment, products are arithmetic strings, quality is the value of that string, and complexity is measured by the length of the string. They find that buyers make more mistakes and that prices are higher when products are more complex. Their experiment differs from ours in that buyers face strong time pressure and are given no information about the objectives

¹⁴ Montero and Sheth (2019) extend this voluntary disclosure design to consider the impact of consultation among receivers and real-world framing, and Sheth (2019) extends this design to consider the impact of competition on voluntary disclosure.

and incentives of sellers, so it is difficult to know what buyers believe about why sellers present products in a complex way. Martin (2015) conducts a lab experiment in which sellers set prices of products of exogenous quality, where like Kalaycı and Potters (2011), products are arithmetic strings, quality is the value of that string, and complexity is measured by the length of the string. He finds that behavior is largely in line with the equilibrium of a model where buyers are rationally inattentive to quality. Gu and Wenzel (2015) present a lab experiment where competing sellers choose whether to “obfuscate” or not, which mechanically changes the degree of naivete in computer buyers, and study how often sellers choose to obfuscation under different policy interventions. More recently, de Clippel and Rozen (2020) implement a lab experiment where senders can obfuscate by requiring receivers to count the number of blue and red balls in a display (the state is given by the larger number of balls). They also find that a majority of senders strategically obfuscate, but this is an equilibrium response in their game, as the desired strategies of senders are only implemented with some probability.

Our work is also related to lab experiments that study vagueness and ambiguity as a way to shroud information. For instance, Serra-Garcia et al. (2011) allow cheap-talk communication to take the form of vague messages (reporting a subset of states). They 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 (natural language) and ambiguous (interval) messages and find that an announcer in coordination games might want to use such messages.

Relative to this literature, we consider complexity as another way to shroud information, under the constraint that the reported information must convey the true state no matter whether it is simple or complex. In this way, it relates to other disclosure experiments where senders are mandated to disclose truthfully but have access to vague messages (reporting a subset of states that contains the true state) which are more complex than precise messages (Deversi, Ispano, and Schwardmann 2019; Li and Schipper 2019). In Li and Schipper (2019), lack of sophistication in disclosure is modelled with finite levels of reasoning using iterated admissibility, which produces the result that the lowest type sends the vaguest (and hence most complex) possible message in the hope a naive receiver will guess a higher type.

Finally, our work relates to a large experimental literature on overconfidence, which is one of the best documented behavioral biases in the lab (see Moore and Healy 2008 for both comprehensive evidence and a review). There is mounting evidence from lab experiments that ego-utility is an important driver of overconfidence about ability and can lead to asymmetric updating in beliefs after receiving feedback about performance (Eil and Rao 2011; Mobius, Niederle, Niehaus, and Rosenblat 2011).

3. Experimental Design

In this section, we present the game of complex disclosure that we implement in the lab. In order to isolate the forces of interest, the game we use takes a simple form. It is based on the sender-receiver framework used to study cheap talk by Cai and Wang (2006) and voluntary disclosure by Jin et al. (2018). We extend this framework to require senders to truthfully disclose the state and allow them to choose the complexity of their messages. After presenting our game, then we give other details about the experimental design.

3.1 The Complex Disclosure Game

In each round, subjects were paired together, and in each pairing, one subject was randomly assigned to be the sender and the other to be the receiver (with equal likelihood). To reduce framing effects, the sender was referred to as the “A Player”, and the receiver was referred to as the “B Player”.

In each round and for each pair, the computer drew a whole number from 1 to 10, called the “secret” number. Each of these numbers was equally likely to be drawn, and both senders and receivers were made aware of this probability distribution.

Each sender was shown the secret number for their pairing and then made their decision about report complexity while the receivers waited. In our main sessions, the sender chose a “report length”, which was a whole number c between 1 and 20. The computer program then randomly selected c integers between -10 and 10 until those

numbers added up to the true state b . Both senders and receivers were told how these numbers were generated.

After all senders made their decisions, the receivers' screens became active. If a sender decided on a report of length c , the receiver they were paired with was shown this message: "The number I received is", followed by a table of c integers ranging from -10 to 10 that added up to the secret number. The instructions explicitly stated that the sender only chooses the report length c and that the table of random numbers is generated by the computer. In the appendix, we present the full instructions and an example of a report with maximum length ($c=20$).

Below the area for the sender's message, receivers were asked to make a guess a of the secret number b , and this guess could be any integer between 1 and 10. The receiver had 60 seconds to view the sender's report and make a guess. If nothing was guessed after that time, a random guess is entered for the receiver. In our main sessions, less than 4% of receivers hit this time limit.

Receiver payoffs, denominated in "Experimental Currency Units" (ECU), were $\frac{1}{2}(b - a)^2$, where b is the secret number and a is the receiver's guess.¹⁵ These payoffs decrease monotonically as the guess moves further from the secret number. The sender payoffs in each round were $\frac{1}{2}(a - b)^2$. These payoffs are independent of the secret number and increase monotonically with receiver guesses because guesses cannot be higher than 10. These payoffs are similar to the quadratic specification found in Crawford and Sobel (1982) when there is a large bias towards higher actions. Because we use just a small number of states and actions, the payoffs could be shown in a table, so that subjects did not need to know or interpret these functional forms.

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

¹⁵ 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.

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

3.2 Experimental Sessions

Our 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. Subjects do not have to be Harvard University students, so we restricted subjects to be no older than 25 years old in order for the subject pool to be more comparable with existing studies that recruit undergraduate students. The software used to run the experiments was the z-Tree software package (Fischbacher 2007).

Each session consisted of 30 rounds of the disclosure game. In each round, subjects were randomly matched into pairs. Each subject could be matched with any other subject in the session and was equally likely to be paired with any given subject. 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%.

The purpose of having subjects play both roles is to ensure that both sides have a good sense for the incentives and actions available to the other side. This design feature might serve to increase strategic sophistication, as prior voluntary disclosure research (e.g. Jin et al. 2018) that has found evidence that playing both roles appears to increase the extent of learning. Thus, if roles had been fixed instead, we might have observed even more evidence of naivete in receivers. However, we do not see much change in behavior over rounds in our experiment, so repeated experience in other role does not seem to have much of an effect on behavior.

At the end of each session, subjects were privately paid in cash a show up fee of \$5 plus all additional earnings they accumulate over the course of the session. ECU were converted to U.S. dollars at a rate of 150 to 1 (rounded up to the nearest dollar). While it is possible for subjects to end up with a negative balance of ECU, because subjects are

paid for every round, this outcome is extremely unlikely and never came close to occurring in the sessions we ran. However, because subjects are paid for every round, there is the potential for intentional variation in play (a “portfolio” strategy), but we find little evidence of such behavior.

3.3 Feedback, Beliefs, Math Test, and Demographics

Our main sessions provide round-by-round feedback. Subjects were told four pieces of information after each round: 1) the actual secret number; 2) the report length chosen by the sender; 3) the receiver’s guess of the secret number, and; 4) their own payoff. After all subjects pressed the “OK” button on the screen containing this feedback, the next round began. To reduce social considerations, subjects in the feedback treatment were not told the payoff 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.

Once all rounds are completed, subjects were asked questions about their beliefs of how other subjects played in their session. First, subjects were asked to guess the average report length that senders chose for each secret number. Second, subjects were asked to guess the average secret number when the sender chose complexity levels between 1 and 5, between 6 and 10, between 11 and 15, and between 16 and 20. The purpose of these questions was to assess whether subject beliefs about sender strategies influenced their decisions as receivers. These guesses were not incentivized, which introduces the possibility of additional noise. However, we find that even with this additional noise, elicited beliefs are fairly accurate for most subjects.

In some sessions, subjects were asked to complete a four-question math test after answering the two belief questions. For each question in this test, subjects were asked to add up 20 numbers, and were paid \$4 if a randomly selected question was correctly answered. Subjects were told that the numbers would sum up to an integer between 1 and 10, that all integers were equally likely, and that the 20 numbers would be generated in the same fashion as in the disclosure game. After completing the math test, subjects answered two additional belief questions. First, they were asked to guess the number of

questions on the math test (from 0 to 4) that they thought they answered correctly. Second, they were asked to guess the average number of questions they thought others answered correctly.¹⁶ These belief questions were also not incentivized.

At the end of the experiment, subjects were asked to complete a questionnaire that includes questions about demographic details. Specifically, subjects are asked for their gender, if they are a native English speaker, their year in school, and if they have a friend participating in that session.

3.4 Robustness Sessions

For robustness, we adopted two alternative treatments. The first alternative replaces round-by-round feedback with “no feedback,” where 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. The no-feedback treatment is designed to contrast with the feedback treatment, so that we can determine whether round-by-round feedback is crucial in driving convergence towards unraveling as in Jin, Luca, and Martin (2018). The second alternative treatment also limits sender choice of report length to the two extremes (c is only 1 or 20) rather than the full range from 1 to 20. The reason for this alternative treatment is to determine whether play is substantially different if the “strategic complexity” of the game is reduced for both senders and receivers. The results of these robustness sessions are provided in Section A.4 of the appendix.

4 Experimental Results

In this section, we first report the results from our main sessions and then explore the possible reasons behind sender and receiver mistakes. For receivers, we estimate a structural model to predict choices both with and without behavioral biases.

¹⁶ The exact wording of the questions was “For yourself, what do you think was the number of rounds (between 0 and 4) answered correctly?” and “For all participants, what do you think was the average number of rounds (between 0 and 4) answered correctly?”

4.1 Results from the Main Sessions

Table 1 summarizes the characteristics for the subjects in our 29 main sessions. In total, we have 294 subjects, all of whom experience both roles (sender and receiver) and receive round-by-round feedback for 30 rounds. Roughly 41% of the subjects are male, 72% are undergraduate students, 85% are native English speakers, and 14% report that they have a friend in the same session.¹⁷ These demographic distributions are similar to the ones reported by Jin, Luca, and Martin (2018), who also conducted experiments in the CLER lab.

4.1.1 Summary of Behavior and Mistakes

Table 2A summarizes sender choice of complexity by secret number. In contrast to the unraveling prediction, the average choice of complexity is 9.728 and increases almost monotonically as the secret number gets smaller.¹⁸ For the two smallest secret numbers (1 and 2), a majority of senders choose the maximum complexity (report length $c=20$) and over 72% choose high complexity ($c \geq 15$). For the two highest secret numbers (9 and 10), a majority of senders choose the simplest report ($c=1$) and over 72% choose low complexity ($c \leq 5$).

Figure 1A depicts the distribution of complexity choices for each secret number, where the size of the bubble represents the number of senders choosing a specific complexity level conditional on a specific secret number. Most senders concentrate on high complexity when the secret number is below 5, and switch to low complexity when the secret number is above 5. If the secret number is exactly 5, sender choices are dispersed across all levels of complexity. These patterns continue in the second half of the rounds, as shown in Figure 1B.

¹⁷ One subject did not report any demographics, and three subjects skipped the question about whether they were native English speakers. Despite these missing values, we include all subjects in the analysis because we can account for these missing values in our subsequent regression analysis.

¹⁸ In a regression of complexity choice onto secret number with individual fixed effects and robust standard errors, the coefficient is negative (-1.496) and statistically significant ($p < 0.001$).

Figure 1A. Frequency and average sender choice of complexity by secret number with 95% confidence intervals (main sessions)

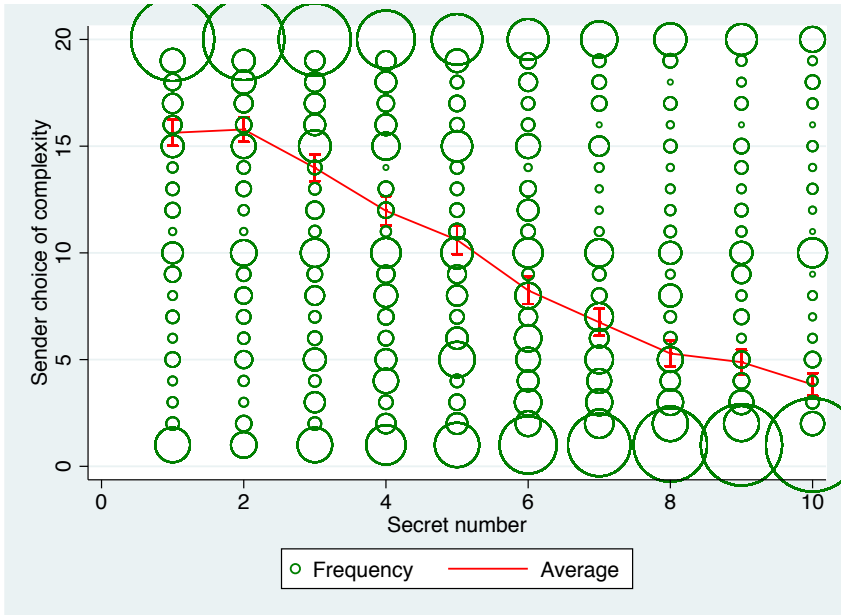
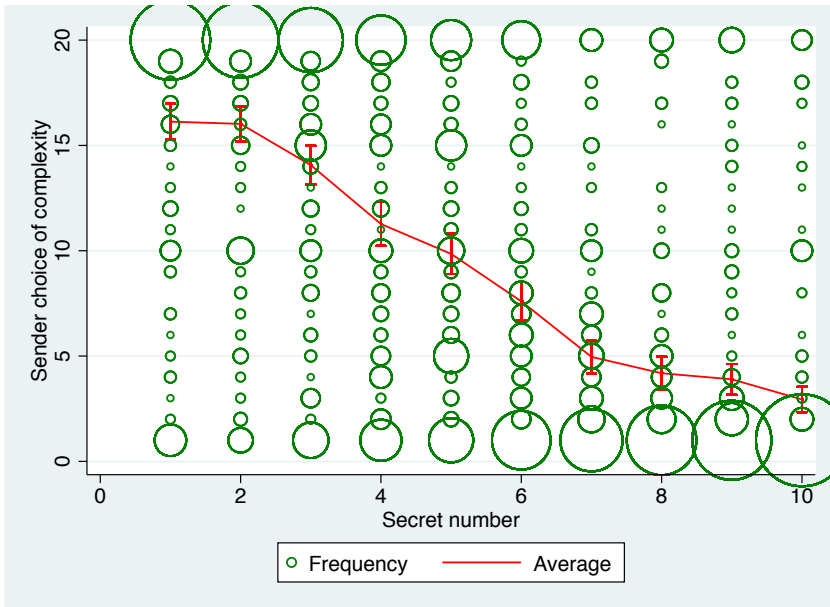


Figure 1B. Frequency and average sender choice of complexity by secret number in the second half of rounds with 95% confidence intervals (main sessions)



Turning to receivers, Table 2B shows that the median receiver guess is correct for every secret number, but the average guess is significantly different from the secret

number for every secret number except for 6 and 7 (using a two-sided t-test and a significance level of 5%). On average, the bias in receiver mistakes reveals much greater over-guessing for low secret numbers (1.183 for secret numbers of 1 and 0.936 for 2) than under-guessing for high secret numbers (-0.367 for 9 and -0.403 for 10).

To further explore receiver behavior, we define the size of receiver mistakes as the absolute distance between the receiver guess and the secret number. As shown in Table 2B, the average receiver mistake size is the highest for the lowest secret number ($c=1$) and decreases almost monotonically with secret number.¹⁹ This is consistent with the fact that senders present simpler reports for higher secret numbers, which reduces the potential for math errors, shortens the response time for receivers, and lowers the probability of receivers not making a decision within the 60 second time limit. For the less than 4% of receivers that are over the time limit, the computer generates a random guess, which can lead to large mistakes. Excluding these observations, receiver mistakes remain large for the smallest secret numbers (0.946 for secret numbers of 1 and 0.777 for 2) as compared to the mistake sizes for large secret numbers (between 0.370 and 0.379 for secret numbers between 6 and 10). In fact, mistake sizes are significantly different between secret numbers of 1 and 10 using a two-sided Wilcoxon rank-sum test (p -value <0.001).

Because receivers observe the complexity of sender reports, Table 3 tabulates how receiver guesses and mistakes vary by the complexity level of sender reports, as well as the secret numbers behind these reports. On average, we observe a small amount of under-guessing for complexity up to a length of 4. For complexity between 5 and 12, the number of observations is smaller, and the average guess fluctuates between over-guess and under-guess. Once complexity is over 12, we observe consistent over-guessing that peaks at the highest level of complexity (0.655 for length 20). Interestingly, the size of receiver mistakes is less monotonic, but is clearly much higher for high complexity than for low and medium complexity.²⁰ These results are robust to excluding rounds where receivers did not make their decision within the time limit. Without those rounds, the

¹⁹ In a regression of mistake size onto secret number with individual fixed effects and robust standard errors, the coefficient is negative (-0.085) and statistically significant ($p<0.001$).

²⁰ In a regression of mistake size onto complexity with individual fixed effects and robust standard errors, the coefficient is positive (0.054) and statistically significant ($p<0.001$).

magnitude of receiver mistakes is slightly lower for the two highest complexity levels (0.783 versus 0.748 for length 19, and 1.284 versus 1.008 for length 20).

Absent behavioral factors, one would imagine that receiver mistakes should be zero for the simplest reports because such report reveals the secret number exactly. In contrast, the data shows an average mistake of 0.243 for length 1 and 0.257 for length 2, which are significantly different from 0 using a two-sided t-test.

To show the joint impact of secret numbers and complexity, Table 4 cross-tabulates secret numbers by low (≤ 5), medium (6-14), and high (≥ 15) levels of complexity. When the secret number is presented simply, receivers tend to have largest mistakes (0.6 on average) for the lowest secret number (1). This is consistent with social preferences because a simple report of a low states is helpful for receivers but harmful for senders, so some receivers may be willing to reciprocate to “honest” senders by sacrificing their own payoff to reward this behavior. However, it may also reflect confusion of about the game. These possibilities are discussed and analyzed further in Section 4.3.

4.1.2 Payoff Losses

So far, we have documented that sender choices of complexity deviate from the unraveling prediction and that receiver guesses deviate from the true state. But do these deviations lead to payoff losses?

To address this possibility, 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.²¹ To do this, we construct the average opponent strategy from our data, determine the expected payoffs from taking each possible action, and then calculate how far the expected payoff for the taken action is from the highest expected payoff.²² For senders, the possible actions are grouped as

²¹ This does not tell us if receiver mistakes are optimal given the noise generated by complexity. The question of whether receivers are acting optimally will be addressed later in the analysis.

²² Because the minimum possible payoff can be negative, we normalize payoffs by subtracting the minimum possible payoff (for the realized state) to the payoffs from taking any action in that state.

low (1-5), medium (6-14), and high (15-20) complexity.²³ For receivers, the possible actions are limited to the guesses available to them, which are integers between 1 and 10.

All of 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 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.²⁴

Table 5 reports the monetary losses that result from actions taken in our main sessions. On average, senders are 15.3% away from the highest expected payoff if they take the empirical distribution of receiver guesses for each complexity group as given. This percentage differs substantially across secret numbers: for the highest secret number (10), sender choice (mostly low complexity) is close to optimal (3.8% loss); but for the lowest secret number, sender choice (mostly high complexity) is still 51.6% away from the highest payoff. This is driven mostly by the failure to always use high complexity when facing a secret number of 1.

We also calculate expected payoffs relative to the payoff that senders would get in the unraveling equilibrium. Because the unraveling equilibrium predicts different receiver behavior that we observe, sender payoffs in equilibrium could be higher or lower than the sender payoff observed in our data. It turns out that sender choice of complexity results in a 71.4% expected gain for secret numbers of 2 and a 3.9% expected loss for secret numbers of 10, relative to the unraveling equilibrium. We cannot do the same exercise for a secret number 1 because the normalized equilibrium payoff is 0.

Table 5 also reports the percentage off from the highest expected payoff that receivers could have achieved if they guessed based just on the observed complexity level (given the empirical distribution of sender types for that complexity level). This

²³ We grouped these actions because some complexity levels are rarely chosen by senders for some secret numbers, thus we could have a non-reliable density in the empirical distribution of sender choice of complexity conditional on these secret numbers. Our results are robust to small changes in the boundaries of these groups, such as having “low” just be lengths of 1 and “high” be lengths of 20.

²⁴ These assumptions may not hold in a dynamic environment that features learning. We will present evidence of learning in Section 4.1.3 and control for these dynamic effects in the regression analyses presented in Section 4.1.4.

deviation is 13.8% for low complexity, 16% for medium complexity, and 16.7% for high complexity. On average, receiver payoffs are 30 to 33% worse than the payoff that receivers would get in an unraveling equilibrium, because receivers would know every state perfectly in this equilibrium. Note that the departure from highest expected payoffs is not readily comparable between senders and receivers, because their payoffs differ in both scale and range.

In short, there are non-trivial sender mistakes and receiver mistakes, even when we measure them in the payoff space. We test the robustness of these results to dynamic effects in Sections 4.1.3 and 4.1.4 and explore the reasons behind these mistakes in Sections 4.2 and 4.3.

4.1.3 Evidence of Learning

To provide detail on sender complexity use over rounds, the first panel of Table 6 also presents how sender payoffs depart from the highest expected payoff over rounds (taking the empirical distribution of receiver behavior as given and fixed over rounds). Overall, we see a gradual improvement from the first 10 rounds (15.9% departure) to the last 10 rounds of the game (14.2%).

Breaking this down by secret number, the biggest improvement comes from the lowest secret number (1), where the departure from the highest payoff drops from 55.1% in the first 10 rounds to 51.5% and 48.4% in the second and third blocks of 10 rounds. Strikingly, this improvement is accompanied by senders *increasing* their choice of complexity for this secret number. In comparison, at the highest secret number (10) senders get closer to the highest payoff (from 5.3% departure in the first 10 rounds to 3.1% in the last 10 rounds) while *decreasing* their use of complexity (from 5.829 to 2.512). The increase in complexity use for low secret numbers and the decrease in complexity use for high secret numbers can be seen by comparing Figures 1A and 1B.

For other secret numbers above 5, we also see senders decrease their use of complexity over the experiment. However, for secret numbers at or below 5, senders continue to use substantial amounts of complex disclosure throughout the experiment, as reflected in an average complexity choice above 10 in the last block of rounds.

Table 6 provides evidence of learning on the receiver side. Throughout the game, the average receiver mistake drops for all three groups of complexity, but the biggest drop occurs for high complexity. Departure from the highest payoff improves as well, while the magnitude of improvement tends to be much larger for medium and high complexity (from ~18% to ~13%) than for low complexity (from 15.3% to 13.7%).

4.1.4 Regression Results

Table 7A presents the results of our regressions based on sender behavior, and Table 7B presents those based on receiver behavior. The motivation for these regressions is to replicate our results while controlling for round-by-round changes in sender and receiver behavior, which we reported in the previous section.

For senders, the dependent variables are sender choice of complexity and the payoff departure from the highest expected payoff. In the first and third columns of Table 7A, we include subject demographics and session fixed effects. Taking a secret number of 1 as the default, Table 7A shows that senders choose significantly less complexity and depart less from the highest payoff when their secret number increases. This is consistent with our results without subject, session, or round controls.

To capture sender learning, we include the round number (1-30) and the interaction with whether the secret number is in the medium (4-6) or high range (7-10). These coefficients suggest that senders learn to *increase* complexity for low states (1-3) but *decrease* complexity for medium and high states. We also include a dummy for the first five rounds, in case the initial learning about the game creates a level effect in choice of complexity. There is little evidence for a difference when controlling for other factors.

Columns (2) and (4) include sender fixed effects, which absorb individual demographics. Results for most coefficients are similar to what we have without individual fixed effects, suggesting that sender choice and learning are not driven by unobserved individual characteristics. For example, this suggests that complexity use is not higher for lower secret numbers just because senders who are more likely to use secret numbers (e.g., unkind senders who enjoy annoying receivers) were by chance assigned lower secret numbers.

Turning to receivers, Table 7B attempts to understand the absolute size of receiver mistakes and receiver's payoff losses when controlling for time trends. Because we want to study the mistakes that receivers actively made, we focus our analysis on the 96% of receiver guesses that are made before the time limit. Since receivers observe the sender's choice of complexity, we include a separate dummy for each complexity level. We control for the same subject-level variables as in the sender regression, but we also include the receiver's response time.

Compared with the default complexity (1), Table 7B shows that receiver mistakes drop significantly for some low complexity levels but increase significantly for some levels of high complexity. This pattern is similar with and without subject fixed effects. Results on payoff losses are less consistent, but once we control for subject fixed effects, payoff losses increase significantly between the default complexity and some high complexity levels. Receivers do appear to lower their guess for high complexity over time (after we control for subject fixed effects). As a result, they depart less from the highest expected payoff with high complexity.

4.2 Reasons Behind Sender Mistakes

From a policy perspective, sender mistakes often capture less interest than receiver mistakes, partly because senders tend to be firms in most field applications, and firms have more resources to overcome their mistakes. However, because subjects play both roles in our experiment, we can hope to learn something about the sources of receiver mistakes by looking at the sources of sender mistakes.

The largest sender losses come from two types of mistakes: using high complexity when the state is high and using low complexity when the state is low. In our main sessions, the former decreases from 11.4% of high-state decisions in the first half of rounds to 7.8% in the second, and the latter occurs in 16.0% of low-state decisions in both the first and second half of rounds.

Both types of sender mistakes could be driven by incorrect beliefs about receiver actions, random errors, or confusion about game form.²⁵ These factors could be ameliorated with experience and feedback, so we might expect their impact to lessen over rounds. However, only the incidence rate of the first mistake – choosing a high complexity level in a high state – decreases over rounds. Evidence that these mistakes might be driven by errors or confusion can be found by comparing the choices a subject makes as a sender and a receiver: there is a positive correlation (0.1344) between the likelihood of a subject choosing a high complexity level in high states as a sender and incorrectly guessing by more than one integer with a simple report when a receiver.

Both types of sender mistakes could also be driven by social preferences. Spite could drive senders to use high complexity when it is not justified in their own payoff, and social norms could drive senders to use low complexity when it is not justified by payoffs. We find some evidence that choosing low complexity in a low state is driven by social preferences by once again comparing the choices a subject makes when they are a sender versus a receiver. If some subjects think that the socially correct action is to disclose simply for even low states, then they might act in this way and reward senders who do the same. In fact, there is a positive correlation (0.2666) between the likelihood of a subject choosing a low complexity level in low states as a sender and over-guessing the state by one integer with a simple report as a receiver.

Because subjects play in both roles, we will include these two possible reasons for sender mistakes – confusion and social preferences – into our baseline model of receiver guesses. However, we find that neither appears to be a major driver of receiver mistakes.

4.3 Reasons Behind Receiver Mistakes

In this section, we study the reasons for the mistakes that receivers make when the secret number is presented in a complex way. Along the way, we also explore the reasons behind the mistakes made with simple reports, but our primary focus is on complex reports because the vast majority of receiver mistakes occur when the secret number is

²⁵ Martin and Munoz-Rodriguez (2019) find evidence of inattention to game form in experiments that use the BDM mechanism.

disclosed with high complexity, and it is these mistakes that justify the complexity that is observed in our experiment. As a consequence, in the subsequent analyses we only use receiver guesses from rounds where senders chose high complexity and where receivers made a guess before the time limit.

We start by modeling receiver mistakes using Logit choice (as in the Quantal Response Equilibrium approach of McKelvey and Palfrey 1995), which assumes that receivers have Logit demand for each action based on the expected payoffs to taking each action given the empirical distribution of opponent actions. This approach has a free parameter often interpreted as the sensitivity of errors to expected payoffs, which we estimate using maximum likelihood. As can be seen in Table 9, the predictions based on this estimated parameter produce an average likelihood of -1.733 and do a reasonably good job at predicting the rates of over-guessing and under-guessing in the experiment. In particular, it is able to capture over-guessing for middle secret numbers.

However, while Logit choice is successful at explaining receiver mistakes, it does not indicate why receivers are making these particular mistakes. The two primary forces we consider for receiver over-guessing of complexity reports are naivete and overconfidence about ability, but we also consider several other possibilities, such as pure boundary effects, social preferences, confusion, and risk preferences.²⁶ To help identify naivete and overconfidence, we elicited subject beliefs about the strategic implications of complex disclosure and their performance in a short math test.

4.3.1 Beliefs about Senders and Math Ability

As mentioned previously, after all 30 rounds of the game were completed, we asked subjects to report what they think the secret number was on average in their session when the report complexity was 1-5, 6-10, 11-15, and 16-20.²⁷ We refer to a subject's guess of the average secret number when the report complexity was 16-20 as their

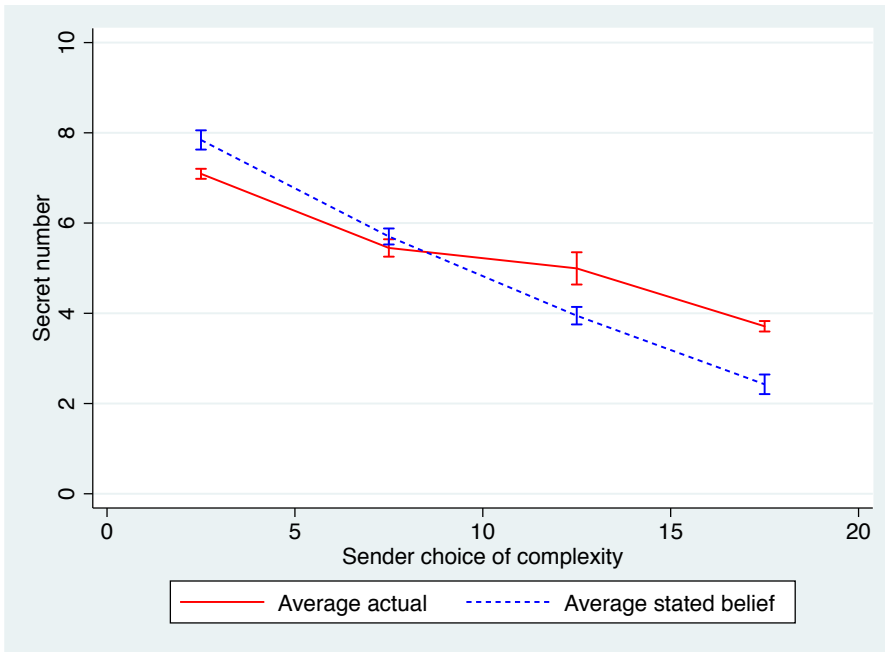
²⁶ Other possible explanations are considered in Section A.5 of the appendix.

²⁷ This belief question uses a different grouping of complexity levels due to a lack of perfect foresight about the clustering of sender actions. Throughout the rest of the paper, we group complexity into low (1-5), medium (6-14) and high (15-20) levels because the empirical distribution of sender choice has much higher density at the two ends (1 and 20) and there is clear bunching at 1, 5, 10, 15 and 20. This difference does not affect our analysis, as we report the summary statistics of stated beliefs separately from other variables.

“complex guess”, and we classify subjects as being “naive” if their complex guess is higher than the actual average secret number when complexity was 16-20 in their session. Across all 294 subjects, 12.6% are classified as naive. When naive, the average amount of naivete is 3.491, which is 98.9% above the actual average secret number in their session.²⁸ Out of the 160 subjects who also completed the math test, 9.4% are classified as naive.

The average answers are presented, along with the actual average secret number for each complexity level, in Figure 2 and in the top panel of Table 8. Given that just 12.6% of subjects are classified as naive, it is unsurprising that the average complex guess (2.510) is lower than both the actual average secret number for such reports (3.626) and the average guess in the game for such reports (4.191).

Figure 2: Average secret number and stated beliefs of average secret number by complexity of 1-5, 6-10, 11-15, or 16-20 with 95% confidence intervals (main sessions)



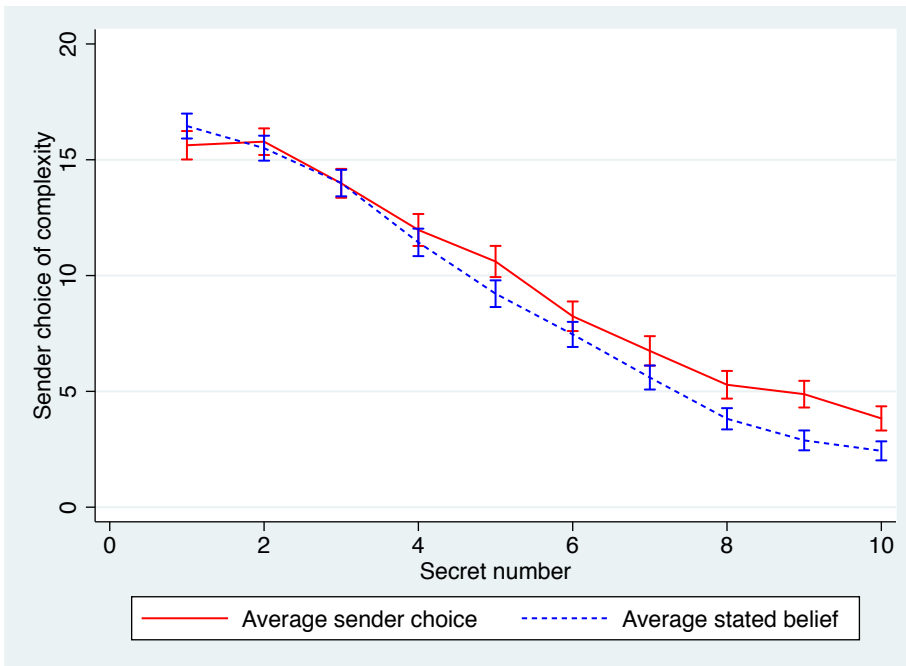
Of course, it is always possible that many more subjects are naive and that the reported beliefs of subjects are not the beliefs used by subjects to play the game (a

²⁸ This is calculated as the difference between complex guess and the actual average divided by the actual average (when the complex guess is higher than the actual average).

possibility raised by Costa-Gomes and Weizsäcker 2008). Also, because we elicit beliefs at the end of the experiment, it could be that subjects are overweighting their experience in the final rounds. However, we observe similar rates of over-guessing in the final 10 rounds, which suggests that the reasons behind receiver mistakes persist into the final rounds.

We also asked subjects to guess the average sender choice of complexity for each secret number, and the average responses are provided in Figure 3. Assuming receivers use these stated beliefs as their prior beliefs at the beginning of each round and only use the observed complexity level (not the content of each report) to determine the value of the secret number, we can infer what they should have guessed via Bayes' Rule. As shown in bottom panel of Table 8, this value (referred to as the "inferred guess") is on average 2.546 for high complexity ($c \geq 15$), which is also lower than the average actual guess in the game for such reports (4.222).

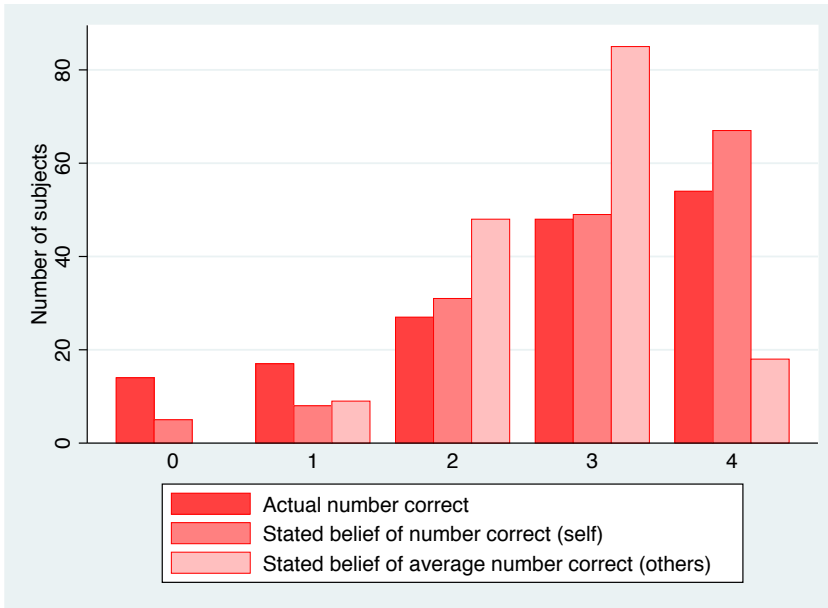
Figure 3: Average sender choice of complexity and stated beliefs of average sender choice of complexity by secret number with 95% confidence intervals (main sessions)



After these strategic beliefs were elicited, 160 subjects were asked to complete a math test that consisted of four questions. Each question required them to sum 20

numbers in a table similar to the most complex table in our game, and these questions were incentivized for correct answers. As shown in Figure 4, 54 subjects got all four questions correct (33.75%), 48 get one wrong (30%), 27 get two wrong (16.88%), and the remaining get either three wrong (10.62%) or all wrong (8.75%).

Figure 4: Math test performance and stated beliefs of math test performance



For the 160 subjects that completed the math test, we classify 33.8% as overconfident because they think they answered more questions correctly than they actually did. When overconfident, the average amount of overconfidence is 1.53. For the overconfident subjects who answered at least one correctly, this is 76.43% over their actual math test performance on average.²⁹

Out of all subjects who completed the math test, 41.88% believe they got all four questions correct and 72.5% believe they got three or four correct (also shown in Figure 4). Both of these rates are higher than the actual fraction of subjects who got this many correct (33.75% and 63.75% respectively). When asked to predict the average number of

²⁹ This is calculated as the difference between their guess of number answered correctly and the actual number they answered correctly divided by the actual number they answered correctly (when their guess of number answered correctly is larger than the actual number answered correctly).

questions that other subjects answered correctly, the average prediction was 2.694, which is close to actual average (also 2.694).

4.3.2 Regressions of Receiver Mistakes on Beliefs

Based on regressions of receiver mistakes onto the extent of naivete and overconfidence (presented in Table 10), we find evidence that when receivers are more naive about complexity and when they are more overconfident about their math performance, receivers over-guess by a larger amount. The relationship between these measures and the size of mistakes is also statistically significant, negative, and of a similar magnitude.

These regressions strongly suggest that both overconfidence and naivete are related to receiver mistakes. However, they do not tell us how much of receiver mistakes are explained by these forces, in part because it does not account for boundary effects. To get a more precise answer, we develop and estimate a structural model.

4.3.3 Structural Model

Because receivers face an involved decision problem, we investigate the sources of receiver mistakes using a partial equilibrium structural model of receiver decision-making when reports are complex. We use a partial equilibrium model instead of a full equilibrium model because senders are largely best responding to receiver behavior. Thus, if we can find a partial equilibrium model that explains receiver behavior, then the behavior in our experiment can be explained using a full equilibrium model with rational senders and possibly irrational (overconfident or naive) receivers.

In our estimations of this model, we pool all receiver guesses for complexity levels above 15 to prove sufficient power for our analysis, but our results are robust to just looking at complexity choices of 20. In addition, we estimate parameters by pooling the choices of all receivers. This is necessary because we have insufficient power to study each individual in isolation. As a consequence, we treat the parameter estimates for this model as coming from a representative agent.

In the model, we assume that a receiver facing a complex message ($c \geq 15$) has prior beliefs about the likelihood of each secret number b given by F (the distribution of sender types using this type of complexity). The receiver then observes a noisy signal of the secret number, which can be interpreted as either an error in summing the numbers or partial attention to the grid of numbers.³⁰ We assume that this noise signal is generated by adding to the secret number an error term e drawn from the distribution G (for this complexity level), so that

$$x = b + e, \text{ where } e \sim G$$

We assume that the distribution of additive errors G has support over the integers $\{-9, -8, \dots, 0, \dots, 8, 9\}$. To increase power, we assume this distribution is symmetric, so the distribution has just 10 parameters to estimate.³¹

Based on the signal x and their prior beliefs F , the receiver forms posterior beliefs γ and takes an action a (makes the guess) that maximizes their expected utility subject to some probability of making strategic errors. This decision rule is given by the following optimization problem:

$$\max_{a \in A} \sum_{b \in B} \gamma(b|x) U_R(a, b)$$

$$\text{where } \gamma(b|x) = \frac{F(b)G(x-b)}{\sum_{b' \in B} F(b')G(x-b')}$$

We also assume that strategic confusion results in a receiver sometimes guessing in a uniform random way. In the Level-k model, this is often designated as the “Level-0” behavior. Because we are using a representative agent model, this is as if some fraction of agents are Level-0 agents. Formally, this means for some fraction of choices, the receiver chooses every action with equal probability.

As a robustness check, we assume that the receiver sometimes uses social preferences that take the form proposed by Fehr and Schmidt (1999). Note that only one parameter of this model (advantageous inequality β) will have bite. For these choices, the decision rule is instead given by the following optimization problem:

³⁰ We consider the possibility of the receiver choosing to get a signal in Section A.6 of the appendix.

³¹ The assumption of symmetry appears to be justified by the data, as we see largely symmetric errors in a math test where there is high complexity but no strategic considerations (see Figure A1 in the appendix).

$$\max_{a \in A} \sum_{b \in B} \gamma(b|x) [U_R(a, b) - \beta(U_R(a, b) - U_S(a, b))]$$

As an additional robustness check, we assume that utility takes the CRRA form, which means that we allow a free parameter α . In this check, we assume that the utility of the receiver is instead given by:

$$\frac{U_R(a, b)^{1-\alpha}}{1-\alpha}$$

We also consider two possible behavioral factors: naivete and overconfidence. We add naivete to our model by assuming that with some probability receivers think that all states are equally likely. In the Level-k approach, this often constitutes Level-1 beliefs: that opponents are guessing randomly. Formally, this means that $\gamma(b|x) = \frac{1}{|A|}$. We add overconfidence to our model by assuming that with some probability receivers think that noise is actually draw from the distribution G' when is actually drawn from G .

4.3.4 Estimating Math Errors

We assume that math error determines the precision of the signal x , and therefore affects the receiver's posterior beliefs about the secret number. We could impose strong assumptions on the distribution of math errors and try to identify it using receiver decisions in the game, but we choose instead to estimate it out-of-sample for cleaner identification. In particular, we estimate the distribution of math errors non-parametrically, using the math errors found in the math test completed after playing the game. The questions in this test have a similar level of complexity as a report with high complexity, but there should be minimal strategic or social considerations when answering these questions, and the payoff function is such that the receiver should report their modal belief of the secret number, regardless of their risk preferences.

By assuming that receivers guess their signal, we can identify from guesses and secret numbers the frequency with which each signal is realized. To estimate G in this way, we used the math test answers for the 160 subjects who completed the math test.³²

³² Because we do not observe guesses when subjects hit the time limit, we exclude these decisions from the estimation.

The resulting estimate places a large mass (72.5%) on no noise ($e=0$), and the average parameter is 4.8 percentage points from the corresponding parameter in a distribution that places all weight on no noise. Our estimate of G is presented visually in Figure A1 in the appendix.

4.3.5 Estimating Strategic Confusion and Social Preferences

We estimate the degree of strategic confusion and the social preferences of the subjects jointly, using the guesses of receivers when the message has been reported in a low complexity ($c \leq 5$). Again, we deliberately use out-of-sample estimation, in order to shy away from confounding factors such as math error. In doing so, we assume that there are minimal interactions between complexity and strategic confusion or social preferences. In practice, it is likely that social considerations when messages are complex are different from when messages are simple, as receivers may feel some positive reciprocity when simple reports are made.

The parameters of this model were estimated using the Nelder–Mead method, and the standard errors were computed using 1,000 bootstrapping samples. The estimates were a 7.4% probability of uniform random choice (with a standard error of 0.007), a 2.3% probability of using social preferences (with a standard error of 0.005), and a 0.658 advantageous inequality parameter (with a standard error of 0.194).

4.3.6 Baseline Predictions

In our baseline model receivers hold correct prior beliefs over the distribution of states given a complex report (equal to the empirical frequency in the main sessions); make math errors in accordance with the estimated distribution G ; understand that their errors come from this distribution, update their beliefs according to Bayes' rule; and then maximize risk-neutral expected utility given their posterior beliefs, but with the estimated probability of strategic confusion (random guessing). Importantly, all of the parameters in this model are estimated out-of-sample.

Even with correct beliefs, this model predicts over-guessing and under-guessing of the extremes because the boundary pushes math errors and strategic errors into the middle of distribution, which then pushes guesses into the middle of the distribution. However, it does not do so symmetrically. Because senders are much more likely to have low secret numbers when they use complexity, receivers should take this into account when they guess, given their uncertainty about the state.

This asymmetry is reflected in the predictions from the model, which are provided in Table 9 along with the predictions given by several variants of this model. For 100,000 simulated draws from the distribution of noise parameters, it predicts over-guessing of 0.712 for low states, -0.181 for middle states, and -1.662 for high states (with an overall average log-likelihood of -1.553). The actual rates of over-guessing were 0.772, 0.096, and -0.891. Because of the strong impact of prior beliefs, the baseline model failed to capture over-guessing for middle states and over-estimated the degree of under-guessing at high states.

A natural robustness check is adding social preferences to the model. Specifically, we add the rate and degree of social preferences estimated out-of-sample, though this is likely to be an overestimate of the actual social preferences for senders who use complex disclosures. For 100,000 simulations, the amended model predicts over-guessing of 0.792 for low states, -0.130 for middle states, and -1.640 for high states. While the model comes closer to predicting the actual rates of over-guessing, the improvements in predicting these rates are small, and the model still fails to capture over-guessing for middle states. In addition, the overall average log-likelihood of -1.547 is only a bit better for the amended model.

Because receivers face uncertainty about secret number, another natural robustness check is adding risk aversion to the model. To estimate this parameter, we conduct a search over a grid of 1,000 values between 0 and 1 using again 100,000 simulations, and the standard errors were computed using 1,000 bootstrapping samples. The parameter that maximizes log-likelihood is set-identified, and the lower bound is 0.010 and the upper bound is 0.135.³³ As Table 9 shows, adding risk aversion to the

³³ The risk aversion parameter is set identified because changes in the parameter value lead to discontinuous changes in the choice probabilities.

baseline model does not noticeably improve the overall average log-likelihood or the predictions of over-guessing.

4.3.7 Behavioral Factors: Naivete and Overconfidence

In our model, naivete is represented in our model as having Level-1 beliefs: believing that all states are equally likely when messages are complex (e.g., that complexity conveys no bad news at all). Receivers with these beliefs would guess that the average secret number was 5.5, which is higher than it actually was, so they would be classified as naive. As reported previously, we classified 12.6% of receivers in our experiment as naive because they guessed that the average secret number was higher than it actually was when messages were complex. We do not know what these receivers believe about the probability of each state when messages were complex, so for simplicity, we assume that all hold Level-1 beliefs. With the assumption that 12.6% of receivers are “Level-1” naive, the model fits the data better than the baseline model, but this assumption does not change receiver decisions enough to produce over-guessing at middle states. Based on simulations of 100,000 decisions, the overall average log-likelihood increases from -1.552 to -1.519. The amended model (still with no free parameters) now predicts over-guessing of 0.733 for low states, -0.146 for middle states, and -1.554 for high states, where the actual rates of over-guessing were 0.772, 0.096, and -0.891.

To determine the degree of overconfidence to add to our model, we compare beliefs about performance on the math test to actual performance on the math test. As reported previously, while 72.5% think performed well, only 63.8% actually performed well. Using this estimate, we amend the baseline model to assume that receivers think they have a 72.5% chance of performing well at math test. In other words, the representative agent believes that there is a 72.5% chance that the error came from G' , which is estimated non-parametrically from the math test using the answers of subjects who actually performed well at the math test. This distribution is shown in Figure A2 in the appendix.

For the model with overconfidence, the predictions for over-guessing are 0.749 for low states, 0.018 for middle states, and -0.904 for high states, where the actual rates of over-guessing were 0.772, 0.096, and -0.891. Thus, the model with overconfidence is able to capture over-guessing for middle states that the baseline model does not capture. For example, the baseline model predicts under-guessing when the secret number is 5 because the distribution of states is heavily skewed towards lower secret numbers, which a fully rational receiver would take into account when guessing. However, the model with overconfidence predicts over-guessing when the secret number is 5 because overconfident receivers largely ignore the heavy skew towards lower secret numbers when guessing.

Based on simulations of 100,000 decisions, the average log-likelihood of the baseline model when overconfidence is added rises from -1.552 to -1.272, which is even higher than the log-likelihood of -1.5194 from the model that includes a portion of Level-1 receivers. We also run a model that includes both Level-1 naivete and overconfidence, and the overall average log-likelihood changes by only a little. In fact, it falls slightly to -1.274. For this combined model, the predictions of over-guessing are 0.761 for low states, 0.027 for middle states, and -0.890 for high states, where the actual rates of over-guessing were 0.772, 0.096, and -0.891.

We also consider an alternative method for estimating overconfidence, which is inspired by the approach for determining distortions of Bayes' rule used in Grether (1980) and Holt and Smith (2009). Our approach, which has a free parameter, is to assume that when updating beliefs, the probability that a signal is observed in a certain state is raised to the power of the parameter. When this parameter is equal to 1, the receiver updates Bayes' rule in the standard fashion. When this parameter is greater than 1, if a signal is more likely in a state (such as the probability of receiving a signal of 7 when the true state is 7), then more weight is given to the state given this signal (such as the probability the true state is 7 given a signal of 7). To estimate this parameter, we conducted a search over a grid of 1,000 values between 1 and 30 using 100,000 simulations, and the standard errors were computed using 1,000 bootstrapping samples. The parameter value that maximizes log-likelihood is also set-identified, but the lower bound is 16.760, which is far from the Bayesian value of 1.

This approach also does a good job at explaining receiver guesses. The overall average log-likelihood is -1.261, which is a bit better than the log-likelihood of -1.272 from the initial model with overconfidence (though this new model does not have free parameter). The predictions for over-guessing are 0.776 for low states, 0.050 for middle states, and -0.800 for high states, where the actual rates of over-guessing were 0.772, 0.096, and -0.891.

4.3.9 Overconfidence and Beliefs

To investigate whether mistakes due to overconfidence can persist even if subjects have correct strategic beliefs, we ran an additional math test using where the likelihood of each answer was skewed exactly as in the experiment. After completing the standard math test, in which each answer was equally likely, 70 subjects faced an additional math test in which the likelihood of each answer followed the empirical distribution of states for complex reports from that session. Subjects were told the distribution, but not that it was taken from the experiment.

For these 70 subjects, the average bias in mistakes for complex reports during the experiment was 1.33 for secret numbers of 1 (n=84) and 1.27 for secret numbers of 2 (n=88). In the additional math test, in which subjects faced the same distribution of states but were told the distribution, the corresponding averages were 0.95 (n=64) and 0.82 (n=62). While admittedly underpowered, a two-tailed t-test does not provide a difference between the average bias at either secret number (p-values of 0.3224 and 0.2631). Thus, we do not have evidence that mistakes would change much in the face of correct beliefs about the underlying distribution.

5. Conclusion and Policy Implications

Our results highlight the incentives for firms to strategically complexify information disclosed to consumers, potentially harming consumers and undermining the effectiveness of disclosure. In our experiment, senders use complex disclosure frequently.

Most of this obfuscation is profitable because receivers make systematic mistakes in assessing complex reports.

The patterns we observe have policy implications as well. For example, many obfuscation theories assume naivete in (a fraction of) consumers, hence consumer education that reduces naivete should alleviate the seller's incentives to obfuscate. But sophistication does not save them from obfuscation if they are overconfident about their ability to comprehend complex reports. Policy tools that target overconfidence can be different from education efforts that target consumer naivete.

Given this, our results suggest that overconfidence might be worth exploring in follow-up lab experiments, field studies on complex disclosure, and theoretical models of complex disclosure. In particular, it would be important to establish for which types of complexity such an effect might exist. For example, whether overconfidence also impacts settings with linguistic complexity or vague messages. If this previously over-looked force is found to matter in those settings also, it might be worth addressing in regulatory policy on complex disclosure, as it has a distinct set of policy implications. For example, overconfidence could produce asymmetry in belief updating, which would make feedback less effective at reducing receiver mistakes, as we observe in our experiment.

The link between overconfidence and complex disclosure could be strengthened in a number of ways. For instance, an involved math test could be used to measure both math errors and overconfidence precisely at an individual level, which could make it possible to generate structural estimates at an individual level. Natural variation in math errors and overconfidence across subjects would then produce a strong test of model fitness. Additional belief elicitation could also be used to enhance reduced-form and structural estimates. For example, beliefs could be elicited about distribution of mistake sizes (e.g., how often off by one, two, etc.) or the subjective probability of number correct (e.g., how likely it is they got 4 correct) to provide a more nuanced sense of overconfidence. The link between overconfidence in the math task and disclosure game could also be more tightly linked by asking receivers about their performance in the disclosure game (e.g., how many guesses were correct in the game). However, there is evidence that overconfident individuals can distort their memories of past performance (Huffman, Raymond, and Shvets, 2018).

Additionally, it might be possible to produce exogenous variation in overconfidence and to examine its impact on the extent of complex disclosure. Examples from the literature include varying the type of future tasks faced (Schwardmann and Van der Weele, 2019), by adding noise into the feedback provided (Fischer and Sliwka, 2018), varying task difficulty (Moore and Healy, 2007), and having overconfident subjects select into the task (Camerer and Lovallo 1999). Also, subjects could be offered, at a cost, access to a calculator that quickly and automatically sums the messages. Overconfident subjects are likely to undervalue such a calculator.

Our results also suggest that a mandate on simplicity can be as important as a mandate on truthful disclosure. More generally, this highlights the potential for regulation aimed at encouraging disclosure to be both simple and salient. The subsequent policymaking challenge becomes identifying which simple and salient disclosures provide the highest welfare (Hershfield and Roese 2015; Caplin and Martin 2020). Another policy implication is seen in sender behavior. Surprisingly, round-by-round feedback does not reduce obfuscation. If anything, learning encourages senders to understand receiver mistakes in low states and exploit it via obfuscation.

A final policy implication is related to disclosure in general. Our results suggest that the unraveling prediction is fragile. Although immediate and repeated feedback can steer voluntary disclosure towards the predictions of unraveling, it fails once we change the setting a little away from simple, voluntary disclosure. How to harvest the benefits of the incentives produced by unraveling remains a challenge in the real world.

References

- Agarwal, S., Chomsisengphet, S., & Liu, C. (2010). The importance of adverse selection in the credit card market: Evidence from randomized trials of credit card solicitations. *Journal of Money, Credit and Banking*, 42(4), 743-754.
- Agarwal, S., Ben-David, I., & Yao, V. (2017). Systematic mistakes in the mortgage market and lack of financial sophistication. *Journal of Financial Economics*, 123(1), 42-58.
- 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.
- Anderson, M. L., Chiswell, K., Peterson, E. D., Tasneem, A., Topping, J., & Califf, R. M. (2015). Compliance with results reporting at Clinical Trials. gov. *New England Journal of Medicine*, 372(11), 1031-1039.
- Armstrong, M., & Chen, Y. (2009). Inattentive consumers and product quality. *Journal of the European Economic Association*, 7(2-3), 411-422.
- Armstrong, M., & Vickers, J. (2012). Consumer protection and contingent charges. *Journal of Economic Literature*, 50(2), 477-93.
- Bederson, B. B., Jin, G. Z., Leslie, P., Quinn, A. J., & Zou, B. (2018). Incomplete disclosure: Evidence of signaling and countersignaling. *American Economic Journal: Microeconomics*, 10(1), 41-66.
- Ben-Shahar, O., & Schneider, C. E. (2014). *More than you wanted to know: The Failure of Mandated Disclosure*. Princeton University Press.
- Bhargava, Saurabh; George Loewenstein and Justin Sydnor (2017). Choose to Lose: Health Plan Choices from a Menu with Dominated Option. *The Quarterly Journal of Economics*, 132(3): 1319–1372.
- Bianchi, M., & Jehiel, P. (2015). Financial reporting and market efficiency with extrapolative investors. *Journal of Economic Theory*, 157, 842-878.
- Blake, T., Moshary, S., Sweeney, K., & Tadelis, S. (2017). Price Salience and Product Choice. Working paper presented at the 2017 NBER Digitization meeting.
- Board, O. (2009). Competition and disclosure. *The Journal of Industrial Economics*, 57(1), 197-213.

- Bollinger, B., Leslie, P., & Sorensen, A. (2011). Calorie posting in chain restaurants. *American Economic Journal: Economic Policy*, 3(1), 91-128.
- 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.
- Camerer, C., & Lovallo, D. (1999). Overconfidence and excess entry: An experimental approach. *American Economic Review*, 89(1), 306-318.
- Caplin, A., & Martin, D. J. (2020). *Framing, Information, and Welfare* (No. w27265). National Bureau of Economic Research.
- Carlin, B. I. (2009). Strategic price complexity in retail financial markets. *Journal of Financial Economics*, 91(3), 278-287.
- Carlin, B. I. (2009). Strategic price complexity in retail financial markets. *Journal of Financial Economics*, 91(3), 278-287.
- Carlin, B. I., Kogan, S., & Lowery, R. (2013). Trading complex assets. *Journal of Finance*, 68(5), 1937-1960.
- Célérier, C., & Vallée, B. (2017). Catering to investors through security design: Headline rate and complexity. *Quarterly Journal of Economics*, 132(3), 1469-1508.
- Chetty, R., Looney, A., & Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, 99(4), 1145-77.
- Costa-Gomes, M. A., & Weizsäcker, G. (2008). Stated beliefs and play in normal-form games. *Review of Economic Studies*, 75(3), 729-762.
- Crawford, V. P., & Iriberri, N. (2007). Level - k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private - Value Auctions? *Econometrica*, 75(6), 1721-1770.
- Crawford, V. P., & Sobel, J. (1982). Strategic information transmission. *Econometrica*, 1431-1451.
- DellaVigna, S. (2018). Structural behavioral economics. National Bureau of Economic Research Working Paper w24797.

- DellaVigna, S., & Pollet, J. M. (2005). Attention, demographics, and the stock market. National Bureau of Economic Research Working Paper w11211.
- DellaVigna, S., & Pollet, J. M. (2009). Investor inattention and Friday earnings announcements. *Journal of Finance*, 64(2), 709-749.
- Deversi, M., Ispano, A., & Schwardmann, P. (2018). Spin doctors: A model and an experimental investigation of vague disclosure. Mimeo.
- Dranove, D., & Jin, G. Z. (2010). Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, 48(4), 935-63.
- Dranove, D., Kessler, D., McClellan, M., & Satterthwaite, M. (2003). Is more information better? The effects of “report cards” on health care providers. *Journal of Political Economy*, 111(3), 555-588.
- Duarte, Fabian and Justine S. Hastings (2012). Fettered Consumers and Sophisticated Firms: Evidence from Mexico’s Privatized Social Security Market. *NBER Working Paper #18582*, accessed at <http://www.nber.org/papers/w18582>.
- Eil, D., & Rao, J. M. (2011). The good news-bad news effect: asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2), 114-38.
- Ellison, G. (2005). A model of add-on pricing. *Quarterly Journal of Economics*, 120(2), 585-637.
- Ellison, G., & Ellison, S. F. (2009). Search, obfuscation, and price elasticities on the internet. *Econometrica*, 77(2), 427-452.
- Ellison, G., & Wolitzky, A. (2012). A search cost model of obfuscation. *The RAND Journal of Economics*, 43(3), 417-441.
- Englmaier, F., Schmöller, A., & Stowasser, T. (2017). Price discontinuities in an online market for used cars. *Management Science*, 64(6), 2754-2766.
- Eyster, E., & Rabin, M. (2005). Cursed equilibrium. *Econometrica*, 73(5), 1623-1672.
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114(3), 817-868.
- Feltovich, N., Harbaugh, R., & To, T. (2002). Too cool for school? Signalling and countersignalling. *RAND Journal of Economics*, 630-649.

- Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171-178.
- Fischer, M., & Sliwka, D. (2018). Confidence in knowledge or confidence in the ability to learn: An experiment on the causal effects of beliefs on motivation. *Games and Economic Behavior*, 111, 122-142.
- Fishman, M. J., & Hagerty, K. M. (2003). Mandatory versus voluntary disclosure in markets with informed and uninformed customers. *Journal of Law, Economics, and Organization*, 19(1), 45-63.
- Fung, A., Graham, M., & Weil, D. (2007). *Full disclosure: The perils and promise of transparency*. Cambridge University Press.
- Gabaix, X., & Laibson, D. (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics*, 121(2), 505-540.
- Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *Quarterly Journal of Economics*, 95(3), 537-557.
- 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. D. (2011). Developing a reputation for reticence. *Journal of Economics & Management Strategy*, 20(1), 225-268.
- Grubb, M. D. (2015). Overconfident consumers in the marketplace. *Journal of Economic Perspectives*, 29(4), 9-36.
- Gu, Y., & Wenzel, T. (2015). Putting on a tight leash and levelling playing field: An experiment in strategic obfuscation and consumer protection. *International Journal of Industrial Organization*, 42, 120-128.
- Hagenbach, J., & Perez-Richet, E. (2018) Communication with evidence in the lab. *Games and Economic Behavior*, 112, 139–165.
- Hanna, R., Mullainathan, S., & Schwartzstein, J. (2014). Learning through noticing: Theory and evidence from a field experiment. *Quarterly Journal of Economics*, 129(3), 1311-1353.

- Heidhues, P., Köszegi, B., & Murooka, T. (2016). Inferior products and profitable deception. *Review of Economic Studies*, 84(1), 323-356.
- Hershfield, H. E., & Roese, N. J. (2015). Dual payoff scenario warnings on credit card statements elicit suboptimal payoff decisions. *Journal of Consumer Psychology*, 25(1), 15-27.
- Hirshleifer, D., & Teoh, S. H. (2003). Limited attention, information disclosure, and financial reporting. *Journal of Accounting and Economics*, 36(1-3), 337-386.
- Holt, C. A., & Smith, A. M. (2009). An update on Bayesian updating. *Journal of Economic Behavior & Organization*, 69(2), 125-134.
- Huffman, D., Raymond, C., & Shvets, J. (2018). Persistent Overconfidence and Biased Memory: Evidence from Managers. Mimeo.
- Jacob, B. A., & Levitt, S. D. (2003). Rotten apples: An investigation of the prevalence and predictors of teacher cheating. *Quarterly Journal of Economics*, 118(3), 843-877.
- Jehiel, P. (2005). Analogy-based expectation equilibrium. *Journal of Economic Theory*, 123(2), 81-104.
- 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. (2018). Is no news (perceived as) bad news? An experimental investigation of information disclosure. National Bureau of Economic Research Working Paper w21099.
- Jovanovic, B. (1982). Truthful disclosure of information. *Bell Journal of Economics*, 36-44.
- Kalaycı, K., & Potters, J. (2011). Buyer confusion and market prices. *International Journal of Industrial Organization*, 29(1), 14-22.
- Lacetera, N., Pope, D. G., & Sydnor, J. R. (2012). Heuristic thinking and limited attention in the car market. *American Economic Review*, 102(5), 2206-36.
- Li, Y. X., and Schipper, B. C. (2018). Strategic Reasoning in Persuasion Games: An Experiment. Mimeo.
- Lowenstein, G., Sunstein, C.R., and Golman, R. (2014). Disclosure: Psychology Changes Everything. *Annual Review of Economics*, 6, 391-419.

- Feng Lu, S. (2012). Multitasking, information disclosure, and product quality: Evidence from nursing homes. *Journal of Economics & Management Strategy*, 21(3), 673-705.
- 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. (2016). No news is good news: Voluntary disclosure in the face of litigation. *The RAND Journal of Economics*, 47(4), 822-856.
- Martin, D. (2015). Rational Inattention in Games: Experimental Evidence. Available at SSRN: <http://ssrn.com/abstract=2674224>.
- Martin, D., & Munoz Rodriguez, E. (2018). Misperceiving Mechanisms: Imperfect Perception and the Failure to Recognize Dominant Strategies. Mimeo.
- Mathios, A. D. (2000). The impact of mandatory disclosure laws on product choices: An analysis of the salad dressing market. *The Journal of Law and Economics*, 43(2), 651-678.
- Matthews, S., & Postlewaite, A. (1985). Quality testing and disclosure. *The RAND Journal of Economics*, 328-340.
- McKelvey, R. D., & Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1), 6-38.
- Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 380-391.
- Mobius, M. M., Niederle, M., Niehaus, P., & Rosenblat, T. S. (2011). Managing self-confidence: Theory and experimental evidence. National Bureau of Economic Research Working Paper w17014.
- Montero, M., & Sheth, J. (2019). Naivety about hidden information: An experimental investigation. CeDEx Discussion Paper Series ISSN 1749 – 3293.
- Sheth, Jesal. (2019). Disclosure of information under competition: An experimental study. CeDEx Discussion Paper Series ISSN 1749 – 3293.
- Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, 115(2), 502.

- Mullainathan, S., Schwartzstein, J., & Shleifer, A. (2008). Coarse thinking and persuasion. *Quarterly Journal of Economics*, 123(2), 577-619
- Perez-Richet, E., & Prady, D. (2011). Complicating to persuade. Mimeo.
- Pope, D. G. (2009). Reacting to rankings: evidence from "America's Best Hospitals". *Journal of Health Economics*, 28(6), 1154-1165.
- Ru, Hong and Antoinette Schoar (2016). Do Credit Card Companies Screen for Behavioral Biases? *NBER Working Paper #22360*, accessed at <http://www.nber.org/papers/w22360>.
- Schwardmann, P., & Van der Weele, J. (2019). Deception and self-deception. *Nature human behaviour*, 3(10), 1055-1061.
- Serra-Garcia, M., van Damme, E., & Potters, J. (2011). Hiding an inconvenient truth: Lies and vagueness. *Games and Economic Behavior*, 73(1), 244-261.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics*, 50(3), 665-690.
- Sitzia, S., & Zizzo, D. J. (2011). Does product complexity matter for competition in experimental retail markets? *Theory and Decision*, 70(1), 65-82.
- Spiegler, R. (2006). Competition over agents with boundedly rational expectations. *Theoretical Economics*, 1(2), 207-231.
- Sullivan, M. (2017). Economic Issues: Economic Analysis of Resort Fees. FTC Economic Issue Paper.
- Trautmann, S. T., & Kuilen, G. (2015). Belief elicitation: A horse race among truth serums. *Economic Journal*, 125(589), 2116-2135.
- Viscusi, W. K. (1978). A note on "lemons" markets with quality certification. *The Bell Journal of Economics*, 277-279.

Table 1. Summary of subject characteristics (main sessions)

Variable	N	Mean	Std. dev.
Number of subjects in the session	294	10.680	2.554
Feedback provided (dummy)	294	1.000	0.000
Male (dummy)	293	0.410	0.493
Undergraduate (dummy)	293	0.720	0.450
Native English speaker (dummy)	290	0.852	0.356
Friend in the session (dummy)	293	0.143	0.351

Note: Observation is per subject. Value is missing if demographic information not provided by the subject.

Table 2A. Summary of sender choices of complexity by secret number (main sessions)

Secret number	N	Sender choice of complexity			High complexity (length \geq 15)	Low complexity (length \leq 5)
		Mean	Median	Std. dev.	Mean	Mean
1	449	15.626	20	6.619	0.728	0.145
2	444	15.782	20	6.157	0.721	0.115
3	464	13.983	17	6.837	0.616	0.19
4	422	11.969	13	7.218	0.486	0.275
5	433	10.607	10	7.13	0.390	0.344
6	453	8.243	6	6.914	0.254	0.455
7	424	6.748	4	6.664	0.198	0.583
8	427	5.286	2	6.288	0.141	0.71
9	447	4.879	1	6.197	0.128	0.729
10	447	3.832	1	5.622	0.094	0.796
Total	4410	9.728	9	7.86	0.378	0.432

Table 2B: Summary of receiver guesses by secret number (main sessions)

Secret number	N	Receiver guess			Receiver mistake bias (guess-truth)	Receiver mistake size (guess-truth)	% of receiver decisions hitting time limit	Conditional on receiver decision before time limit	
		Mean	Median	Std. dev.	Mean	Mean	Mean	Receiver mistake bias (guess-truth)	Receiver mistake size (guess-truth)
1	449	2.183	1	2.326	1.183	1.183	5.57%	0.946	0.946
2	444	2.923	2	2.209	0.936	1.045	8.11%	0.659	0.777
3	464	3.399	3	1.462	0.399	0.601	5.39%	0.328	0.492
4	422	4.232	4	1.458	0.232	0.611	3.08%	0.191	0.538
5	433	5.169	5	1.378	0.169	0.566	3.23%	0.146	0.489
6	453	6.031	6	1.167	0.031	0.446	3.97%	0.051	0.377
7	424	6.887	7	1.234	-0.113	0.424	2.12%	-0.067	0.376
8	427	7.724	8	1.289	-0.276	0.407	1.17%	-0.237	0.37
9	447	8.633	9	1.377	-0.367	0.438	2.01%	-0.311	0.379
10	447	9.597	10	1.574	-0.403	0.403	0.67%	-0.372	0.372
Total	4410	5.663	6	2.885	0.182	0.614	3.56%	0.128	0.509

Table 3. Summary of receiver guess by sender choice of complexity (main sessions)

Complexity	N	All receiver decisions Mean values					Conditional on receiver decision before time limit Mean values		
		Secret number	Receiver guess	Receiver mistake bias (guess-truth)	Receiver mistake size (guess-truth)	% hitting time limit	Response time if before time limit	Receiver mistake bias (guess-truth)	Receiver mistake size (guess-truth)
1	1259	7.504	7.466	-0.038	0.243	0.40%	9.15	-0.038	0.236
2	214	6.967	6.925	-0.042	0.257	0.00%	8.95	-0.042	0.257
3	140	6.429	6.407	-0.021	0.15	0.00%	13.21	-0.021	0.150
4	104	5.962	5.885	-0.077	0.135	0.00%	13.35	-0.077	0.135
5	190	5.600	5.684	0.084	0.179	0.00%	18.15	0.084	0.179
6	91	5.527	5.582	0.055	0.231	1.10%	18.85	0.089	0.200
7	89	5.685	5.629	-0.056	0.146	1.12%	21.5	-0.023	0.114
8	117	5.325	5.299	-0.026	0.402	0.85%	23.67	-0.052	0.379
9	74	4.932	5.068	0.135	0.405	0.00%	25.45	0.135	0.405
10	263	5.54	5.388	-0.152	0.479	1.90%	28.77	-0.143	0.438
11	42	5.500	5.476	-0.024	0.595	2.38%	34.25	0.049	0.537
12	69	4.87	4.783	-0.087	0.841	1.45%	35.54	0.000	0.765
13	54	4.778	5.222	0.444	0.778	0.00%	35.56	0.444	0.778
14	39	4.974	5.513	0.538	0.795	2.56%	37.08	0.632	0.737
15	190	4.384	4.463	0.079	0.753	3.16%	36.55	0.071	0.712
16	71	3.592	4.000	0.408	0.662	7.04%	37.21	0.273	0.424
17	90	4.467	4.789	0.322	1.033	5.56%	40.32	0.306	0.847
18	96	4.292	4.573	0.281	1.01	9.38%	42.45	0.195	0.839
19	115	4.07	4.296	0.226	0.783	6.96%	40.33	0.243	0.748
20	1103	3.455	4.11	0.655	1.284	9.79%	42.76	0.477	1.008
Total	4410	5.482	5.664	0.182	0.614	3.56%	24.93	0.128	0.509

Table 4. Summary of receiver mistake size by secret number and sender choice of complexity (main sessions)

Secret number	All receiver decisions			Conditional on receiver decision before time limit		
	Mean values of receiver mistake size (lguess-truthl)			Mean values of receiver mistake size (lguess-truthl)		
	Low complexity (1-5)	Medium complexity (6-14)	High complexity (15-20)	Low complexity (1-5)	Medium complexity (6-14)	High complexity (15-20)
1	0.6	0.386	1.437	0.6	0.386	1.126
2	0.216	0.795	1.234	0.216	0.795	0.873
3	0.273	0.144	0.846	0.273	0.124	0.691
4	0.198	0.376	0.961	0.198	0.35	0.839
5	0.181	0.426	1	0.162	0.372	0.88
6	0.204	0.432	0.896	0.19	0.392	0.74
7	0.142	0.366	1.321	0.138	0.33	1.18
8	0.228	0.672	0.033	0.228	0.548	0.93
9	0.23	0.641	1.404	0.222	0.603	1.098
10	0.239	0.776	1.357	0.239	0.776	1.077
Total	0.225	0.469	1.133	0.221	0.434	0.91

Table 5: Departure from highest expected payoff (main sessions)

Panel A: Senders		
Secret number	Fraction of payoff loss from highest expected payoff given empirical distribution of opponent behavior	Fraction of payoff loss from payoff in the unraveling equilibrium
1	0.516	*
2	0.320	-0.714
3	0.152	-0.160
4	0.110	-0.043
5	0.103	-0.016
6	0.073	0.006
7	0.077	0.028
8	0.078	0.039
9	0.059	0.041
10	0.038	0.039
Total	0.153	-0.088

Panel B: Receivers		
Complexity	Fraction of payoff loss from highest expected payoff given empirical distribution of opponent behavior	Fraction of payoff loss from payoff in the unraveling equilibrium
Low (1-5)	0.138	0.299
Medium (6-14)	0.160	0.330
High (15-20)	0.167	0.311
Total	0.153	0.308

* In the unraveling equilibrium, senders with a secret number of 1 earn the minimum possible payoff. After normalizing this payoff to 0, it is not possible to calculate the fraction of payoff loss from zero.

Table 6: Summary of dynamics (main sessions)

Panel A		Sender choice of complexity			Fraction of sender payoff loss from highest expected payoff		
		Mean			Mean		
Secret number	Round 1-10	Round 11-20	Round 21-30	Round 1-10	Round 11-20	Round 21-30	
1	14.454	16.461	16.032	0.551	0.515	0.484	
2	15.357	15.993	15.958	0.311	0.322	0.326	
3	13.264	15.026	13.693	0.159	0.146	0.150	
4	12.673	12.679	10.467	0.107	0.125	0.097	
5	11.669	9.878	10.13	0.105	0.105	0.098	
6	9.526	7.646	7.545	0.068	0.084	0.066	
7	9.475	5.719	5.036	0.091	0.079	0.061	
8	6.764	5.218	3.693	0.086	0.081	0.064	
9	6.326	5.455	3.093	0.058	0.068	0.050	
10	5.829	3.5	2.512	0.053	0.031	0.031	
Total	10.624	9.786	8.774	0.159	0.156	0.142	

Panel B		Receiver mistake size (Iguess-truthI)			Conditional on before time limit		
		Mean			Mean		
Complexity	Round 1-10	Round 11-20	Round 21-30	Round 1-10	Round 11-20	Round 21-30	
Low (1-5)	0.254	0.206	0.222	0.247	0.202	0.219	
Medium (6-14)	0.472	0.518	0.410	0.442	0.476	0.375	
High (15-20)	1.274	1.099	1.004	1.015	0.947	0.751	
Total	0.719	0.604	0.520	0.585	0.526	0.418	

Panel C		Fraction of receiver payoff loss from highest expected payoff			Conditional on before time limit		
		Mean			Mean		
Complexity	Round 1-10	Round 11-20	Round 21-30	Round 1-10	Round 11-20	Round 21-30	
Low (1-5)	0.153	0.126	0.137	0.153	0.126	0.137	
Medium (6-14)	0.182	0.148	0.138	0.181	0.148	0.136	
High (15-20)	0.189	0.161	0.147	0.183	0.155	0.130	
Total	0.174	0.143	0.141	0.171	0.141	0.135	

Table 7A: Regressions of sender behavior (main sessions)

	Dependent Variable Complexity		Dependent Variable Payoff departure from the highest	
Secret number=2	0.147 (0.423)	-0.252 (0.418)	-0.197*** (0.0609)	-0.199*** (0.0625)
Secret number=3	-1.467*** (0.476)	-1.297** (0.559)	-0.361*** (0.0714)	-0.360*** (0.0736)
Secret number=4	-1.630** (0.749)	-1.894** (0.711)	-0.418*** (0.0551)	-0.421*** (0.0567)
Secret number=5	-2.884*** (0.689)	-3.358*** (0.610)	-0.426*** (0.0533)	-0.429*** (0.0545)
Secret number=6	-5.226*** (0.753)	-5.405*** (0.633)	-0.452*** (0.0504)	-0.454*** (0.0512)
Secret number=7	-5.485*** (1.019)	-5.742*** (0.994)	-0.433*** (0.0501)	-0.436*** (0.0512)
Secret number=8	-7.134*** (0.886)	-7.393*** (0.858)	-0.438*** (0.0533)	-0.435*** (0.0543)
Secret number=9	-7.363*** (0.952)	-7.614*** (0.937)	-0.453*** (0.0487)	-0.453*** (0.0496)
Secret number=10	-8.386*** (0.842)	-8.278*** (0.928)	-0.476*** (0.0519)	-0.479*** (0.0529)
First 5 rounds	-0.298 (0.297)	-0.387 (0.267)	0.00783 (0.0116)	0.00683 (0.0116)
Round	0.0427* (0.0241)	0.0409 (0.0249)	-0.000657 (0.00117)	-0.000601 (0.00114)
Round * (4<=secret number <=6)	-0.139*** (0.0262)	-0.130*** (0.0212)	0.000730 (0.00117)	0.000685 (0.00109)
Round * (7<=secret number <=10)	-0.216*** (0.0281)	-0.219*** (0.0301)	-0.000183 (0.00118)	-0.000298 (0.00114)
Individual demographics	Yes	No	Yes	No
Individual fixed effects	No	Yes	No	Yes
Observations	4,410	4,410	4,399	4,399
R-squared	0.350	0.529	0.381	0.438

In parentheses are robust standard errors clustered by session. *** p<0.01, ** p<0.05, * p<0.1. In Session 34, receivers' actual play is such that the highest payoff for draw=1 is 0 after our normalization, so we cannot calculate fraction of payoff departure from 0. That is why columns (3) and (4) have 11 less observations. Regressions without individual fixed effects include dummies indicating whether demographics are missing and session fixed effects. Sample includes all sessions with complete demographic information.

Table 7B: Regressions of receiver behavior (main sessions)

	Dependent variable Receiver mistake size (lguess-truth)		Dependent variable Payoff departure from the highest expected payoff	
Sender choice of complexity = 2	-0.0191 (0.0590)	0.0228 (0.0627)	-0.0229* (0.0121)	-0.0126 (0.0125)
Sender choice of complexity = 3	-0.162* (0.0849)	-0.0517 (0.0651)	-0.0177 (0.0152)	-0.0151 (0.0169)
Sender choice of complexity = 4	-0.199** (0.0860)	-0.130 (0.0784)	0.000650 (0.0206)	0.00884 (0.0212)
Sender choice of complexity = 5	-0.185** (0.0777)	-0.169** (0.0793)	0.0168 (0.0160)	0.0213 (0.0172)
Sender choice of complexity = 6	-0.143 (0.116)	-0.111 (0.0862)	0.00231 (0.0216)	0.00419 (0.0193)
Sender choice of complexity = 7	-0.286** (0.109)	-0.0927 (0.113)	0.00188 (0.0201)	0.0110 (0.0231)
Sender choice of complexity = 8	-0.0171 (0.126)	-0.0480 (0.123)	0.0127 (0.0192)	0.0147 (0.0209)
Sender choice of complexity = 9	-0.0755 (0.141)	-0.143 (0.149)	0.0483* (0.0258)	0.0440 (0.0294)
Sender choice of complexity = 10	0.00256 (0.109)	0.000952 (0.0981)	0.0343* (0.0178)	0.0442** (0.0198)
Sender choice of complexity = 11	0.0232 (0.238)	-0.0690 (0.260)	-0.0156 (0.0295)	-0.00468 (0.0288)
Sender choice of complexity = 12	0.334 (0.199)	0.344 (0.209)	0.0118 (0.0204)	0.0214 (0.0256)
Sender choice of complexity = 13	0.319 (0.271)	0.451 (0.278)	0.0559** (0.0259)	0.0608* (0.0303)
Sender choice of complexity = 14	0.319 (0.315)	0.379 (0.308)	0.0776** (0.0305)	0.0870** (0.0359)
Sender choice of complexity = 15	0.341* (0.193)	0.475** (0.226)	0.0212 (0.0186)	0.0386* (0.0221)
Sender choice of complexity = 16	0.0371 (0.194)	0.115 (0.209)	-0.0173 (0.0254)	0.00332 (0.0270)
Sender choice of complexity = 17	0.391 (0.240)	0.570** (0.257)	0.0679** (0.0324)	0.0755* (0.0379)
Sender choice of complexity = 18	0.373 (0.239)	0.408* (0.238)	0.0187 (0.0338)	0.0145 (0.0358)
Sender choice of complexity = 19	0.317 (0.258)	0.274 (0.250)	0.0253 (0.0314)	0.0438 (0.0332)
Sender choice of complexity = 20	0.587*** (0.181)	0.605*** (0.197)	0.0197 (0.0188)	0.0315 (0.0226)
First 5 rounds	-0.0342 (0.0614)	-0.0710 (0.0615)	0.0127 (0.0106)	0.0118 (0.0109)
Round	-0.00335 (0.00349)	-0.00115 (0.00291)	-0.000374 (0.000527)	-0.000328 (0.000534)

Round * Medium complexity (6-14)	-0.00372 (0.00628)	-0.00748 (0.00463)	-0.00138* (0.000752)	-0.00145* (0.000766)
Round * High complexity (15-20)	-0.00705 (0.00480)	-0.0149*** (0.00459)	-0.00151* (0.000789)	-0.00198** (0.000859)
Response time (in seconds)	0.00824* (0.00411)	0.0114*** (0.00376)	0.000482* (0.000281)	0.000397 (0.000333)
Individual demographics	Yes	No	Yes	No
Individual fixed effects	No	Yes	No	Yes
Observations	4,253	4,253	4,253	4,253
R-squared	0.094	0.279	0.040	0.127

All regressions are conditional on receivers making a guess within the 60 second time limit. In parentheses are robust standard errors clustered by session. *** p<0.01, **p<0.05, *p<0.1.

Regressions without individual fixed effects include dummies indicating whether demographics are missing and session fixed effects.

Table 8: Summary of receiver guess and stated beliefs (main sessions)

Panel A: Complex guess (stated belief of average secret number for a given complexity)						
Complexity	All received decisions			Conditional on before time limit		
	Secret number	Receiver guess	Complex guess	Secret number	Receiver guess	Complex guess
	Mean	Mean	Mean	Mean	Mean	Mean
1-5	7.091	7.064	7.813	7.091	7.064	7.823
6-10	5.448	5.396	5.756	5.447	5.404	5.756
11-15	4.701	4.835	3.867	4.655	4.818	3.848
16-20	3.626	4.191	2.51	3.636	4.055	2.471

Panel B: Inferred guess (secret number inferred from stated beliefs of sender choices)						
Complexity	All receiver decisions			Conditional on before time limit		
	Secret number	Receiver guess	Inferred guess	Secret number	Receiver guess	Inferred guess
	Mean	Mean	Mean	Mean	Mean	Mean
Low (1-5)	7.091	7.064	7.845	7.091	7.064	7.849
Medium (6-14)	5.338	5.344	4.893	5.326	5.354	4.891
High (15-20)	3.712	4.222	2.546	3.72	4.097	2.526

Note: Out of all receiver decisions, 6.8% have a missing value for inferred guess because those subjects indicate that senders will never choose some complexity level.

Table 9: Summary of structural estimation of receiver guesses of high complexity reports before time limit (main sessions)

Variable	Actual	Logit	Baseline	Social Preferences	Risk Aversion	Naivete	Over-confidence	Over-Confidence + Naivete	Over-weighting
Mean log-likelihood		-1.733	-1.553	-1.547	-1.553	-1.519	-1.272	-1.274	-1.261
Total log-likelihood		-2641	-2366	-2357	-2366	-2316	-1939	-1941	-1921
Parameter (lower)		0.047			0.010				16.760
Std. error		0.066			0.206				0.347
Parameter (upper)					0.135				23.076
Std. error					0.196				1.500
Receiver bias (guess-truth)									
Mean values									
Secret number									
1-3	0.772	0.707	0.712	0.792	0.712	0.733	0.749	0.761	0.776
4-7	0.096	0.038	-0.181	-0.130	-0.181	-0.146	0.018	0.027	0.050
8-10	-0.891	-0.641	-1.662	-1.640	-1.662	-1.554	-0.904	-0.890	-0.800
Average distance		0.125	0.369	0.332	0.369	0.315	0.038	0.027	0.142
Receiver bias (guess-truth)									
Mean values									
Secret number									
1	1.126	1.100	0.983	1.069	0.983	1.005	1.001	1.018	1.034
2	0.711	0.632	0.706	0.790	0.706	0.726	0.735	0.744	0.742
3	0.431	0.335	0.407	0.476	0.407	0.428	0.475	0.485	0.517
4	0.249	0.154	0.193	0.258	0.193	0.214	0.271	0.281	0.289
5	0.222	0.044	-0.082	-0.033	-0.082	-0.059	0.033	0.039	0.068
6	0.040	-0.044	-0.287	-0.255	-0.287	-0.262	-0.141	-0.133	-0.103
7	-0.462	-0.154	-1.172	-1.128	-1.172	-1.067	-0.432	-0.421	-0.385
8	-0.684	-0.335	-1.368	-1.341	-1.368	-1.264	-0.636	-0.625	-0.548
9	-0.980	-0.632	-1.725	-1.699	-1.725	-1.614	-0.926	-0.911	-0.840
10	-1.077	-1.100	-2.009	-1.999	-2.009	-1.898	-1.268	-1.250	-1.115
Average distance		0.159	0.393	0.370	0.393	0.340	0.091	0.093	0.936

Table 10: Regressions of receiver over-guessing in complex rounds if completed math test (main sessions)

	Dependent variable: Receiver mistake (guess-truth)		Dependent variable: Receiver mistake size (lguess-truthl)	
Sender choice of complexity = 15	-0.274*	-0.233	-0.0920	-0.0559
	(0.143)	(0.178)	(0.175)	(0.221)
Sender choice of complexity = 16	0.223	0.243	-0.212	-0.0716
	(0.248)	(0.325)	(0.268)	(0.376)
Sender choice of complexity = 17	-0.473*	-0.481	-0.113	-0.155
	(0.266)	(0.365)	(0.112)	(0.128)
Sender choice of complexity = 18	-0.297	-0.647**	-0.294	-0.338
	(0.233)	(0.298)	(0.193)	(0.288)
Sender choice of complexity = 19	-0.558	-0.356	-0.239	-0.121
	(0.320)	(0.399)	(0.295)	(0.382)
First 5 rounds	-0.313	-0.211	-0.283	-0.247
	(0.184)	(0.187)	(0.173)	(0.205)
Round	-0.00277	-0.00125	-0.0149*	-0.0207**
	(0.00711)	(0.00909)	(0.00722)	(0.00898)
Size of naivete (complex guess-actual average if >0)	0.245**		0.233	
	(0.0829)		(0.139)	
Size of overconfidence (guess correct - actual correct if >0)	0.236**		0.295*	
	(0.108)		(0.144)	
Individual demographics	Yes	No	Yes	No
Individual fixed effects	No	Yes	No	Yes
Observations	813	813	813	813
R-squared	0.085	0.270	0.085	0.359

All regressions are conditional on receivers making a guess within the 60 second time limit. In parentheses are robust standard errors clustered by session. *** p<0.01, ** p<0.05, * p<0.1. Regressions without individual fixed effects include dummies indicating whether demographics are missing and session fixed effects.