Hiding a Flaw?
Experimental Evidence on Multi-Dimensional Information Disclosure*

Agata Farina† Mario Leccese‡

September 14, 2025

Abstract

Economic agents often rely on communication to overcome information asymmetries. When information is verifiable, the unraveling principle predicts full disclosure, as senders should reveal all private information. Yet, experimental evidence shows that with a single relevant dimension, senders frequently withhold unfavorable information, and receivers remain overly optimistic about what is left undisclosed. Many real-world settings, however, involve multiple attributes of interest to receivers. We study how such multi-dimensionality affects disclosure choices and inference. In a controlled laboratory experiment, we show that increasing the number of attributes the sender can disclose from one to two increases the disclosure of unfavorable information. Belief elicitation reveals that this pattern is driven by senders' expectations: they anticipate greater skepticism about undisclosed dimensions, particularly when disclosing only one. Receivers, in contrast, do not become more skeptical, failing to recognize the strategic link between disclosure choices across attributes.

*We are very grateful to Guillaume Fréchette, Ginger Zhe Jin, and Maher Said for several insightful discussions. We also thank Chiara Aina, Jeremy Bertomeu, Joyee Deb, Emel Filiz-Ozbay, Jimena Galindo, Sam Kapon, Pierre Liang, Alessandro Lizzeri, Daniel Martin, Elif Osun, Jacopo Perego, Andrew Schotter, Andrew Sweeting, Victor van Pelt, Ece Yegane, and the participants to the 2025 AES Annual Conference, 2025 DISS Workshop, AES Weekly Webinars, 2022 North America ESA Meeting and to seminars at NYU and University of Maryland for their helpful comments. Mario Leccese gratefully acknowledges the research support provided by the Behavioral and Social Sciences Dean's Research Initiative Award, University of Maryland, College Park. An earlier circulated version was entitled "Hiding a Flaw: An Experimental Analysis on Multi-Dimensional Information Disclosure". The usual disclaimer applies.

†Department of Economics, Princeton University. Email: agatafarina@princeton.edu.

‡Questrom School of Business, Boston University. Email: leccese@bu.edu.

1 Introduction

In many markets, decision-makers must evaluate multiple attributes that are difficult to observe, while another party may hold verifiable information about them. Hiring managers, for instance, need to assess a candidate's problem-solving and communication skills, while investors face similar challenges when judging a firm's financial health and managerial capabilities. When such attributes remain hidden, uninformed decision-makers rely on beliefs, which can distort choices and generate market failures (Stiglitz, 2000).

A central result in information economics—the *unraveling principle*—argues that, when communication frictions are minimal, strategic forces push the informed party (sender, he) to disclose all verifiable information about otherwise unobservable attributes. If information about some attribute is not disclosed, the uninformed side (receiver, she) interprets this omission in the most pessimistic way. More optimistic beliefs cannot be sustained in equilibrium, because any sender with more favorable information would have the incentive to reveal it, requiring a revision of the receiver's inference (Viscusi, 1978; Grossman and Hart, 1980; Grossman, 1981; Milgrom, 1981).

Engaging in this reasoning requires receivers to think hypothetically about the different realizations of the attributes with which the sender may be endowed and the strategic choices connected to them. In particular, the receiver should be able to account for the selection in the information she sees and extract payoff-relevant information from such selection. Failure to make such reasoning can lead to naïveté about undisclosed information and provide the sender with a profitable deviation from the theoretical equilibrium. Jin, Luca and Martin (2021) show that, in a game without immediate feedback about the outcomes of prior disclosure choices, senders endowed with a single attribute (i.e., a uni-dimensional information space) tend to disclose favorable realizations but withhold unfavorable ones, best responding to receivers' lack of skepticism. Receivers then hold incorrect beliefs about the sender's strategy, partly ignoring the link between the sender's private information and disclosure choice.[1]

While these findings highlight how unraveling may fail when the information space is uni-dimensional, in many contexts, senders have private verifiable information about multiple attributes. This makes the strategic nature of the decision to disclose (or not) more salient because the receiver is more likely to observe some

[1] This issue has been widely investigated in the theoretical and behavioral economic literature. For instance, in the context of cursed equilibrium (Eyster and Rabin, 2005; Cohen and Li, 2022; Fong, Lin and Palfrey, 2025), analogy-based expectation equilibrium (Jehiel, 2005; Jehiel and Koessler, 2008), level-k reasoning and cognitive hierarchy model (Nagel, 1995; Camerer, Ho and Chong, 2004), shrouded attributes (Gabaix and Laibson, 2006; Heidhues, Kőszegi and Murooka, 2016).

information being disclosed and some other being withheld. For example, in an investor presentation, a firm can decide to voluntarily disclose information on its business performance to attract capital. Suppose a firm sells its product in Boston and New York City, but in its presentation, it only discloses information about its performance in New York City. How might investors interpret this omission? Would they infer that performance in Boston is worse than in New York City? Would their inference change if the firm only operated in Boston? If so, could such differences in investors' beliefs about non-disclosed operations influence the firm's disclosure strategy?

In this paper, we use a lab experiment to study whether and to what extent increasing the dimensionality of the information space, thereby enabling partial disclosure, affects the receiver's skepticism about non-disclosed attributes and hence senders' disclosure incentives. If simultaneously observing different disclosure decisions can help the receiver form beliefs about the sender's strategic choices, then a larger information space should improve the receiver's ability to perform the hypothetical reasoning required to trigger the sender's disclosure. For example, if one dimension is disclosed and the other is not, it should be easier for the receiver to infer that the hidden information is less favorable to the sender, suggesting the link between the private information and the disclosure choice. Anticipating this dynamic, the sender's disclosure rate might be higher when the information space is larger.

This mechanism and the treatment we propose to test it are inspired by the recent experimental literature studying the failure of contingent thinking. Several studies show that subjects' performance in tasks requiring the assessment of multiple contingencies improves when the problem is presented in a way that helps them focus on all relevant contingencies (see Niederle and Vespa, 2023 for a survey). Similarly, in our context, observing the disclosure choice over multiple attributes may serve as an aid to the receiver, helping her to consider different contingencies and uncover the true strategic meaning behind hidden states.

Specifically, we consider two games between two players and design an experiment that builds on Jin, Luca and Martin (2021). In the control (C), the first stage has the sender privately observing a number drawn from $\{1, 2, 3, 4, 5\}$—which can be thought of as the quality of the unique attribute characterizing a firm's product or a job candidate's ability—and choosing whether or not to reveal it. If the sender chooses to disclose, then the actual number is shown to the receiver, meaning that dishonest reporting is prohibited (i.e., information is verifiable). After observing the number disclosed, or nothing if the sender chooses not to disclose, the receiver guesses the number. Payoffs are designed in a way such that the sender wants the receiver to guess as high as possible, while the receiver wants to guess the number as accurately as possible. Our main treatment (T) modifies the control game by adding an attribute to the sender's type. Two numbers are independently drawn from $\{1,2,3,4,5\} \times \{1,2,3,4,5\}$, the sender chooses what to disclose (both numbers, one number, or nothing) and the receiver makes a guess for each attribute. The theoretical prediction in our control treatment is information unraveling: the sender should fully disclose the private information and the receiver should be fully skeptical about undisclosed numbers (Milgrom, 1981). In our treatment, unraveling could fail, but partial disclosure of information (i.e. disclosure of a single number) can only be consistent with the withholding of the worst possible number. Overall, from a theoretical standpoint, the introduction of the second dimension leads to equal or less information transmission. However, we conjecture that increasing the number of attributes from one (as in C) to two (as in T) raises the sender's disclosure rate, defined as the probability that any secret number is disclosed. The intuition relies on the role of the second number as an aid to the receiver: observing potentially different choices within the two numbers might make the strategic incentives behind disclosure decisions more salient, triggering the receivers' skepticism and making senders more pessimistic about the profitability of undisclosed information. This dynamic would lead the empirical results closer to full information transmission.

Our results show that by increasing the number of attributes that a sender can disclose from one to two, the disclosure rate of unfavorable information significantly increases, whereas it does not change for the favorable one. In particular, the aggregate disclosure rate of 1 and 2 significantly increases from 16.5% to 33.4%, leading to a growth in the overall disclosure rate from 57.2% to 65%. However, in both treatments receivers' guesses upon no-disclosure are close to the unconditional average number (i.e., 3). This suggests that providing senders with the opportunity to engage in partial disclosure by increasing the dimensionality of the information space—and hence making information withholding more salient to receivers—fails to improve receivers' ability to account for the selection in disclosed numbers.

To gain a better understanding of the mechanism driving these results, we elicit players' beliefs using an incentivized task. We use the binarized scoring rule to elicit the full distribution of two random variables, i.e., receivers' beliefs over secret numbers and senders' beliefs over receivers' guesses.[2] Differently from the previous literature, we directly elicit beliefs over payoff-relevant distributions, which do not necessarily coincide with the strategy of the other player. This difference in the elicitation procedure allows us to directly

[2]Incentive compatibility of the binarized scoring rule is independent of a subject's risk attitude (Allen, 1987; McKelvey and Page, 1990; Schlag and Van der Weele, 2013; Hossain and Okui, 2013).

identify what the receivers' beliefs over undisclosed information are, rather than deriving them through the Bayes rule, taking as given the beliefs over the sender's strategy. Since several experimental studies documented that subjects' belief updating tends to be inconsistent with the Bayes rule (e.g., Benjamin, 2019), our methodology allows for more accurate estimates of receivers' beliefs.

By analyzing subjects' beliefs, we find that when the number of attributes increases, senders believe that receivers will become more skeptical about non-disclosed states and disclose more frequently. This result sheds light on the channel through which disclosure rates increase following the growth in the dimensionality of the information space: senders believe that the additional dimension will improve receivers' ability to reason through contingencies and partially correct for the lack of skepticism after undisclosed information. On the contrary, receivers do not show any significant change in their beliefs. This result is somewhat surprising: senders interpret the experimental intervention as anticipated and appropriately respond to it, while receivers do not internalize the strategic nature of the different disclosure choices.

To sum up, the dimensionality of the information space and the resulting scope for partial disclosure systematically shape strategic communication. While the unraveling prediction is a central benchmark across markets with verifiable information, existing theoretical and empirical work has largely focused on settings with a single dimension of quality. We demonstrate experimentally that adding a second dimension increases the disclosure of unfavorable information, driven by senders' expectations that receivers will be more skeptical toward undisclosed attributes. Yet receivers' beliefs and actions do not change substantially. This shows that disclosure incentives can be shaped by anticipated audience reactions, even when those reactions never occur. Because settings in which information receivers care about information about multiple attributes of an option are pervasive (e.g., in consumer goods, financial reporting, labor markets, and public services), our findings highlight how disclosure strategies alone can reshape market communication, even in the absence of new regulations or changes in enforcement.

Related Literature. This paper complements a growing empirical literature on voluntary disclosure, whose applications range across accounting, finance, industrial organization, and regulation (for reviews, see: Verrecchia, 2001; Milgrom, 2008; Dranove and Jin, 2010; Beyer et al., 2010.)

The unraveling equilibrium is a central result in information economics (Viscusi, 1978; Grossman and Hart, 1980; Grossman, 1981; Milgrom, 1981).[3] In the last few years, many papers have studied the empirical va-

[3] Other papers have studied more complex settings in which unraveling can fail. In particular, full disclosure is not an equilibrium

lidity of this theoretical prediction and constantly detected the existence of residual information asymmetry after communication. A natural question is what the forces behind the failure of the theoretical result are and whether it is possible to correct the information environment to make communication fully effective. Our work explores the relationship between the dimensionality of the information space, the sender's strategic incentives, and the receiver's inference given the disclosed information.

By testing in the lab a verifiable disclosure game in the spirit of Milgrom (1981), Jin, Luca and Martin (2021) find that unraveling is incomplete: the sender withholds unfavorable evidence due to the receiver's insufficient skepticism about undisclosed information. Other papers in the literature find excessive concealment of unfavorable evidence, suggesting deviations of different magnitudes from the theory. Forsythe, Isaac and Palfrey (1989) and King and Wallin (1991*a*) also show that the failure of complete unraveling is mainly due to the receiver's unsophistication. As in the Jin, Luca and Martin (2021)'s "feedback" treatment, they find that such deviation from the equilibrium tends to disappear as subjects learn.[4] Forsythe, Lundholm and Rietz (1999) finds evidence of information withholding when comparing the effectiveness of voluntary disclosure and cheap talk in reducing adverse selection. While different in their primary focus, other papers suggest the robustness of incomplete unraveling in strategic disclosure games. These studies allow for vague messages (Li and Schipper, 2020; Deversi, Ispano and Schwardmann, 2021), realistic features akin to those typical of field experiments (Montero and Sheth, 2021), unawareness of the evidence available (Li and Schipper, 2018), sender's competition (Sheth, 2021), group disclosure (Avoyan and Onuchic, 2024), complex disclosure (Jin, Luca and Martin, 2022), or partial commitment power (Fréchette, Lizzeri and Perego, 2022).[5] In particular, Sheth (2021) studies a setting in which the receiver observes the disclosure choice of two competing senders, guesses each sender's private information, and chooses with which opponent to interact with higher probability (e.g., which combination of guess and sender's type will be more likely to de-

prediction if the disclosure is costly (Jovanovic, 1982; Verrecchia, 1983) or if receivers are uncertain about the sender's information endowment (Dye, 1985; Jung and Kwon, 1988).

[4] Bourveau, Breuer and Stoumbos (2020) study disclosure behavior in the streetcar industry of the 1890s and find that learning increases disclosure.

[5] Benndorf, Kübler and Normann (2015) studies a disclosure game in which the receiver's role is played by a computer and so cannot exhibit naivëte. The paper still finds failures in complete information unraveling. Dickhaut et al. (2003) and King and Wallin (1991*b*) study a disclosure game in which the sender may be uninformed (as in Dye, 1985 and Jung and Kwon, 1988). Overall, they find results that are in line with the theoretical predictions. Finally, de Clippel and Rozen (2020) studies a communication game in which senders can obfuscate verifiable information and find evidence of strategic obfuscation in the data. Incomplete unraveling has also been studied in the field. For instance, Mathios (2000) and Jin and Leslie (2003) provide evidence of the withholding of information in the contexts of food nutrition labels and restaurants' hygiene grade cards and Brown, Camerer and Lovallo (2012) in the contexts of pre-release movie critics. In a different setting, Bertomeu, Ma and Marinovic (2020) show that managers strategically withhold information from investors.

termine her payoff). This setting relates to ours because the receiver can observe and compare two disclosure choices, but since such choices are taken by different subjects, the inference that the receiver can make about the underlying disclosure strategy of the sender is limited.

Hagenbach and Perez-Richet (2018), Penczynski, Koch and Zhang (2022), and Farina et al. (2024) study settings in which the sender can selectively disclose verifiable information. However, the main focus of our work differs from theirs. Hagenbach and Perez-Richet (2018) investigate the relationship between the outcome of the disclosure problem and different degrees of preference alignments between the players. Penczynski, Koch and Zhang (2022) study the impact of sender competition on the effectiveness of communication when selection of information is possible. Farina et al. (2024) more generally investigate selective disclosure as a force in communication, varying the constraints the sender faces about the amount of evidence that can be disclosed. The paper studies how senders select noisy verifiable evidence about their (unidimensional) type and to what extent receivers account for such selection.[6]

Furthermore, we contribute to the literature by studying how individuals make inferences when presented with strategically selected information. Specifically, we propose that observing multiple strategic choices might increase the salience of the omitted information and countervail agents' inability to account for selection in the disclosed information. We view this variation in the information space as a tool that helps receivers better visualize the strategic nature of the senders' choices – which does not face any disclosure friction/constraint on any of the dimensions – and more effectively engage in the hypothetical reasoning required for their inferences.[7] Our findings suggest that the effectiveness of this aid depends on the nature of the decision problem—for instance, whether it pertains to the sender's or the receiver's task. This provides further evidence of the potentially major role that difficulties with contingent reasoning can play in leading individuals to routinely make sub-optimal choices. In this sense, our paper is related to the recent literature that studies failure in contingent thinking and selection-neglect (Charness and Levin, 2009, Esponda and Vespa, 2014, 2018, 2023; Agranov et al., 2018; Martínez-Marquina, Niederle and Vespa, 2019; Enke, 2020; Araujo,

[6] Brown and Fragiadakis (2019) and Degan, Li and Xie (2023) also study settings in which the sender faces some constraint over the amount of evidence that can be disclosed, and compare a treatment in which evidence is strategically selected to one in which it is selected randomly.

[7] Other papers in the literature have theoretically studied selective disclosure when frictions are present, and derived conditions under which partial disclosure can be, instead, part of senders' optimal behavior. Kirschenheiter (1997) characterizes optimal disclosure choices when a manager can disclose – at a cost – two different types of signals which can separately affect the market value of the firm. Pae (2005), Dye and Finn (2007), Dziuda (2011) and Guttman, Kremer and Skrzypacz (2014) develop different models to show how partial (or even no) disclosure can be part of the equilibrium when receivers are uncertain over the number of signals a sender can disclose.

Wang and Wilson, 2021; Ali et al., 2021; Aina, Amelio and Brütt, 2023; Barron, Huck and Jehiel, 2024).

The rest of the paper is organized as follows. Section 2 presents the theoretical model underlying the experiment, whereas Section 3 describes the details of our experimental design. Section 4 presents the results of the study, and a conclusion is offered in Section 5.

2 Theoretical Framework

We study two variations of a one-shot disclosure game involving two players, a sender and a receiver.

In the first variation, which will serve as the "control" in our empirical analysis, the sender draws her type, $\theta \in \Theta \subset \mathbb{N}$ from a discrete uniform probability distribution with full support over Θ. Next, θ realizes, and the sender observes it, while the receiver only knows the ex-ante distribution. At this point, the sender can take two possible actions: reporting her type or saying nothing. In other words, the set of messages available to the sender when the state is θ is $M(\theta) = \{\theta, \emptyset\}$. After observing the sender's message—which could be empty—the receiver has to take a guess $a \in A = \Theta$. We can interpret this action of the receiver as guessing the sender's type.

The true type and the receiver's guess determine the payoff of both players. In particular, the sender's utility is given by a function $U_S^C(a)$, which is strictly increasing in the receiver's action and is independent of the state. The receiver's utility, instead, is given by the function $U_R^C(a, \theta)$, which is concave in the receiver's action a and is maximized at $a = \theta$. Intuitively, the receiver benefits from guessing numbers that are close to the true type of the sender, while the sender benefits from the guess being as high as possible. The nature of the utility functions determines a conflict of interest between sender and receiver, which gets stronger as the sender's type becomes lower.

The argument in Milgrom (1981) can be used to show that there always exists a sequential equilibrium in which the sender always reports her type (unless it is the lowest element of Θ) and if the message is empty, the receiver takes the action that is optimal given the lowest element of Θ. This implies that when the state is the minimal element of Θ, the sender is indifferent between reporting it or not, so any mixture between the two actions is consistent with equilibrium behavior. We can refer to this equilibrium as the "unraveling equilibrium" or simply "unraveling". We note that in the version of the game that we test in the lab, unraveling is not necessarily the unique sequential equilibrium because the action space is not rich enough.[8] In Section 3 we discuss what are the relevant equilibria in our parametrization and we highlight our main testable predictions.

The second variation, which will be the "treatment" studied in the empirical analysis, considers a similar game with the exception that both the type and the action space are bi-dimensional. The type of the sender is now characterized by a vector of two independently drawn realizations $(\theta_1, \theta_2) \in \Theta^2$. The message space is also richer, since now the sender can not only decide to fully disclose her type or disclose nothing, but she can also decide to reveal only one dimension of her type, i.e., $M(\theta_1, \theta_2) = \{(\emptyset, \emptyset), (\theta_1, \emptyset), (\emptyset, \theta_2), (\theta_1, \theta_2)\}$. After observing the sender's message, the receiver has to take two guesses $(a_1, a_2) \in A^2 = \Theta^2$.

After the message is sent and the guesses are taken, one dimension of the problem, i.e., a pair (θ_i, a_i) for $i = 1, 2$, is randomly drawn, with equal probability, to determine players' payoffs according to the same payoff functions of the control game. This means that $U_S^T(a_1, a_2) = \frac{1}{2}U_S^C(a_1) + \frac{1}{2}U_S^C(a_2)$, while $U_R^T(a_1, a_2, \theta_1, \theta_2) = \frac{1}{2}U_S^C(a_1, \theta_1) + \frac{1}{2}U_S^C(a_2, \theta_2)$. It is worth emphasizing that the structure of the payoff functions, and in particular the fact that $U_R^T(a_1, a_2, \theta_1, \theta_2)$ is separable in (a_1, θ_1) and (a_2, θ_2), ensures that receiver can take the two guesses independently, eliminating any concern that different attitudes toward risk of experimental subjects may drive the empirical results.

The goal of this different game is to allow the receiver to observe two disclosure decisions of the same player at once. Such information may be used to make inferences about the strategy used by the opponent and affect the naïvete in receivers' behavior documented by the previous literature.

As in the control game, the argument used in Milgrom (1981) guarantees that there always exists a sequential equilibrium in which the sender always reports both dimensions of her type (unless they are equal to the lowest type of Θ) and if at least one dimension of the message is empty, the receiver makes a guess equal to the lowest value of Θ for that dimension. This implies that when one θ_i is equal to the lowest element of Θ, the sender is indifferent between reporting it or not. This is the only case in which, in such an equilibrium, we can observe partial disclosure. Again, we will refer to this equilibrium as the "unraveling" equilibrium. Similarly to before, this sequential equilibrium is not unique but we show that partial disclosure never arises as an equilibrium outcome. In Section 3 we describe in more detail the set of relevant equilibria

[8] The more parsimonious message space allows us to implement a more tractable belief elicitation task, which we will detail in Section 3. Nonetheless, the fact that our results for the control game are similar to Jin, Luca and Martin (2021), where unraveling is the unique sequential equilibrium, allows us to exclude that our players behave according to a potentially different equilibrium. Appendix B describes the full set of equilibria that can arise in our game.

3 Experimental Design

We conducted the experiment in the Experimental Economics Laboratory at the University of Maryland (EEL-UMD) during the Spring and Fall semesters of 2022. We recruited 120 subjects from the University of Maryland undergraduate student pool via ORSEE (Greiner, 2015). Each subject could participate only in one session. We used the software oTree to design the experiment (Chen, Schonger and Wickens, 2016). Each subject was provided with a personal computer terminal and separated from the others through dividers. We conducted a total of 10 sessions, 4 sessions for the control and 6 sessions for the treatment. Each session included 8 to 16 subjects. We had an equal number of subjects for each treatment (60 subjects). The average session lasted one hour and the average payment was $19 including the $7 show-up fee. During the experiment, subjects earned their payoffs in "Experimental Currency Units" (ECUs), with a conversion rate of 200 ECU for $1. The final payment was presented to the subjects already converted in US dollars, rounded to the nearest non-negative whole dollar amount. Subjects received paper instructions that were read out loud.[9] The experiment was divided into two parts. Subjects were informed of such division at the beginning of the experiment, but the detailed instructions for each part were given separately at the appropriate time to avoid any possible distortion of the subjects' behavior. A summary of the key information was provided on the screen before the beginning of each part of the experiment. At the end of each session, subjects were paid, privately and in cash, their participation fee plus the earnings from the experiment.

In the experiment, subjects needed to complete 50 rounds, either of the control or of the treatment game. The first part of the experiment consisted of 30 rounds in which subjects were asked to play the main game (we will refer to this as Part I of the experiment). The second part of the experiment consisted of the remaining 20 rounds, in which subjects played the main game and the additional belief elicitation task (we will refer to this as Part II of the experiment).

We use a between-subjects design with fixed roles and random matching. At the beginning of the session, subjects were assigned to either the role of sender or to the role of receiver and they kept such role for the whole experiment. The choice to focus on "fixed roles" aims at mimicking real market interactions more closely, and follows the recent experimental literature on verifiable disclosure games (e.g., Sheth, 2021; Montero and Sheth, 2021; Fréchette, Lizzeri and Perego, 2022). In each round, subjects were randomly and

[9]Detailed instructions are reported in Appendix C.

anonymously matched with another subject in the session playing a different role. We did not provide any feedback during the game and subjects were made aware of their final payoff only at the end of the session.[10]

3.1 Main Game: Control and Treatment

Following Jin, Luca and Martin (2021), we refer to the sender as S player and to the receiver as R player. The interactions between pairs of players followed the theoretical model, with $\Theta = \{1,2,3,4,5\}$, and happened without any interruption for the first 30 rounds. Each round in 1–30 proceeded as follows.

After observing her type, each sender selected the message to send to the receiver. Then, subjects playing as receivers were shown the disclosed information and required to make their guess(es). Each guess needed to be a number from the set $A = \{1,2,3,4,5\}$ and there was no time limit for the receiver's choice.

We ran two different kinds of sessions: control sessions, where the sender's type was represented by a single random integer between 1 and 5, and treatment sessions, where a sender's type in any given round was represented by a pair of independently drawn random integers between 1 and 5. In their analysis, Jin, Luca and Martin (2021) show that their qualitative results are robust to an increase in the number of states from 5 to 10, suggesting that a larger type space for the sender does not by itself modify the dynamic of the game.

Given the modified type space, the set of messages a sender could send varied between treatment and control sessions. In control sessions, senders could only choose between disclosing their type (secret number) or disclosing nothing. If the latter was chosen, the receiver received "No message" on their screen. In treatment sessions, senders could choose between disclosing both their secret numbers, only one of them, or none. As in the control sessions, in case the sender decided to not disclose anything, the receiver received "No message" on their screen for both secret numbers. If, instead, only one secret number was disclosed, the receiver saw the disclosed dimension and received "No message" for the other one. For what concerns the receiver's task, in the control sessions, it was required to indicate only one guess, while in the treatment sessions, the number of guesses increased to two (one for each secret number).[11]

For the payoff functions, we follow closely Jin, Luca and Martin (2021), using $U_S^C(a) = 110 - 20|5-a|^{1.4}$ and $U_R^C(a,\theta) = 110 - 20|\theta - a|^{1.4}$. Given the limited number of states and actions, the payoffs were

[10] In Part I, subjects were remunerated for all 30 rounds. In Part II, where beliefs are elicited, subjects were paid for a randomly selected round. The instructions contained information about the payment scheme.

[11] Appendix C includes the screens shown to the senders and the receivers when making their choices.

provided in a table, as shown in Appendix C. Subjects were informed of all these features of the game in the instructions, and the payoff table was present on the screen every time they needed to make a decision.

We did not provide the subjects with any feedback. This makes our control treatment equivalent to the "no feedback-fixed role" treatment in Jin, Luca and Martin (2021). Such a decision is partially motivated by the fact that such treatment is the one in which the failure of unraveling seems more likely to persist.[12] In addition, our main goal is to understand whether a larger information space can facilitate receivers' reasoning about senders' strategic choices, leading to more correct beliefs, and whether this can affect senders' behavior. Not providing feedback makes changes in beliefs the main driver of our results, which turns out to be a desirable property.

3.2 Belief Elicitation

After completing the first 30 rounds, subjects moved to Part II of the experiment. During this second part, in each round, they were asked to complete the same task described above and then to answer a few additional questions. In particular, senders were asked to state the probability that they associated with each guess of a receiver after a given message was shown. Similarly, receivers were asked to state the probability they associated with each secret number after receiving a given message.[13] In this way, we can elicit each sender's belief over receivers' potential guesses as well as each receiver's beliefs over secret numbers for a given disclosure choice. Such belief elicitation also involves cases in which the secret number is disclosed, providing us with a natural way to check subjects' understanding of the task.

We incentivize the belief elicitation task using the binarized scoring rule. Subjects stating beliefs that are close enough to the truth are rewarded with a monetary prize, measured in ECUs. Such a prize allocation rule makes incentive compatible for subjects to state their true beliefs and allows us to retrieve the distribution used by the subjects in their strategic choices. All the details on how the prize is allocated for the two treatments are reported in Appendix D.

We do not provide subjects with details about this mechanism during the instruction period to avoid them engaging in complex computations that might bias the reporting rather than helping in their decision pro-

[12] Jin, Luca and Martin (2021) find evidence of incomplete information unraveling in all their treatments, but providing subjects with feedback increases disclosure across rounds. In particular, disclosure rates almost converge to full disclosure in the "feedback and random role" treatment, where subjects have the greatest chance to learn due to role reversal and continuous feedback.

[13] Appendix C illustrates the belief elicitation tasks completed by subjects in the lab.

cess (Danz, Vesterlund and Wilson, 2022). Instead, we directly tell subjects that if they wish to maximize the likelihood of receiving the monetary prize, it is in their best interest to answer truthfully to the belief elicitation questions.[14]

The questions we ask to elicit beliefs are based on the messages that senders disclosed in that particular round. Given the setting, if the selected message for a given sender-receiver pair was the outcome of the disclosure choice made by the sender in that pair in that particular round, the behavior in the main game might be distorted. In particular, a sender could take a suboptimal disclosure choice to make the subsequent belief elicitation questions easier and, in this way, increase the likelihood of the prize allocation. To prevent this from happening, we ask each pair of subjects questions concerning the disclosure choice of an unknown different pair in that round. This mechanism allows us to retrieve subjects' beliefs about the unknown variables, without distorting their behavior. Indeed, subjects are randomly and anonymously re-matched in every round, so they are not aware of the identity of their opponent. Eliciting beliefs about the behavior of their opponent in a particular round or about the behavior of any other subject playing in the opposite role does not make a difference: the beliefs each subject forms cannot depend on the identity of the subject they are matched with.

How we elicit beliefs differs from both Jin, Luca and Martin (2021) and Sheth (2021). Jin, Luca and Martin (2021) do not incentivize the belief elicitation task, while we remunerate subjects in an attempt to reduce the noise in their answers. Both Jin, Luca and Martin (2021) and Sheth (2021) elicit receivers' beliefs over the frequency of senders who have disclosed or not a particular secret number during the session. However, the distribution that is relevant for receivers when choosing their action is the probability of each secret number given a particular disclosure choice. Given the elicited beliefs, this conditional distribution can only be backed out by applying the Bayes' rule. However, this requires assuming that receivers update their beliefs via the Bayes' rule, which can be viewed as a strong assumption, especially in light of the failure in players' rationality documented in this kind of sender-receiver games. Our approach allows us to directly elicit the receiver's beliefs over secret numbers given a disclosure choice, relaxing the assumption that receivers update their beliefs via the Bayes' rule.

[14]In the instructions, we informed subjects that they could ask for the details on the prize allocation mechanism at the end of the experiment. The document in Appendix D was available and carefully explained if requested. Out of 60 subjects, only one student required more information about the mechanism.

3.3 Testable Hypotheses

The main theoretical prediction for our control game is information unraveling. Despite the discreteness of the action space, full information transmission remains the theoretical benchmark to assess our empirical findings.[15] In this equilibrium, the receiver is extremely skeptical about the undisclosed information and takes action $a = 1$ after no disclosure.

In the treatment game, the equilibrium analysis suggests two possible outcomes. The first is information unraveling, where the receiver takes action $a = 1$ following no disclosure. The second outcome involves the disclosure of $(\theta_1, \theta_2) \notin \{(1,1),(1,2),(2,1),(2,2),(1,3),(3,1)\}$ and the complete withholding of the remaining states. Here, the receiver would choose $(a_1, a_2) = (2,2)$ when both states are withheld and guess 1 for any undisclosed state when the sender engages in partial disclosure. In both these equilibria, partial disclosure does not play a significant role and it can only arise in case the hidden dimension is equal to 1.

These equilibrium predictions, combined with the empirical results obtained by Jin, Luca and Martin (2021) in the context of a similar game with uni-dimensional information space—i.e., our control—constitute the conceptual framework that we use to derive testable hypotheses on the sender's and the receiver's behavior in our experiment.

Starting with the sender, Jin, Luca and Martin (2021) shows the failure of unraveling in our control game. Our focus is on how disclosure rates change in response to an increase in the dimensionality of the information space, which characterizes our treatment game. Although the theory predicts that our treatment should *not* display higher disclosure rates, we conjecture that this may occur in practice. The reason is that a multi-dimensional information space allows for selective disclosure, thus making information withholding more salient and receivers more skeptical about undisclosed states.

Sender - Hypothesis 1: Increasing the dimensionality of the information space leads to higher disclosure rates.

The second main prediction of the theory concerning the sender's behavior is that in the treatment game partial disclosure should never be part of the optimal strategy unless the undisclosed state is $\theta = 1$. This is because, in equilibrium, receivers are extremely skeptical about undisclosed pieces of evidence, disincentivizing such behavior. We note that infrequent use of partial disclosure is not only consistent with the theoreti-

[15] See Appendix B for a detailed explanation of the equilibria selection.

cal equilibrium but also with senders believing that a multi-dimensional information space enables receivers to better assess all possible contingencies.

Sender - Hypothesis 2: Partial disclosure accounts for a negligible share of senders' choices.

To gain a better understanding of the mechanism driving the conjectured increase in disclosure rates, we elicit and analyze the sender's beliefs. In effect, the sender cannot directly observe changes in the receiver's skepticism towards undisclosed states. Rather, he forms beliefs about the receiver's responses to undisclosed information, both in the treatment and the control. Consistently with our first hypothesis, we expect senders to believe that receivers will be more skeptical about an undisclosed state when the state space is multi-dimensional.

Sender - Hypothesis 3: Senders believe that receivers will take lower actions after undisclosed information in the treatment compared to the control.

For what concerns the receiver's behavior, Jin, Luca and Martin (2021) document that receivers tend to be insufficiently skeptical. We examine how the receiver's skepticism is affected by information withholding. We conjecture that by making the sender's strategy more salient, an additional state should improve the receiver's ability to assess the strategic nature of the disclosure choice, especially when the receiver simultaneously observes a piece of information being disclosed and another being withheld (partial disclosure).

Receiver - Hypothesis 1: When the sender withholds information, receivers make lower guesses in the treatment than in the control. This is particularly true in the case of partial disclosure.

Finally, we also elicit receivers' beliefs to study how these are related to the change in guessing behavior. Consistent with the previous hypothesis, receivers should believe that undisclosed states are, on average, lower in the treatment compared to the control, particularly when partial disclosure is used.

Receiver - Hypothesis 2: Receivers believe that the undisclosed state is lower in the treatment compared to the control. This is particularly true in the case of partial disclosure.

4 Results

To analyze the impact of increasing the dimensionality of the information space on the behavior of senders and receivers in our game, we compare the results between treatment and control sessions. In this section,

Table 1: Disclosure Rates across Treatments

	One-number (C)			Two-numbers (T)		
	Obs	Mean	Std. Dev.	Obs	Mean	Std. Dev.
$\theta = 1$	337	0.059^{*}	0.237	610	0.157^{*}	0.364
$\theta = 2$	282	0.291^{***}	0.455	607	0.512^{***}	0.500
$\theta = 3$	313	0.725	0.447	624	0.769	0.421
$\theta = 4$	288	0.917	0.277	567	0.907	0.291
$\theta = 5$	280	0.946	0.226	592	0.929	0.257
Average	1,500	0.572^{**}	0.495	3,000	0.650^{**}	0.477

Notes: We test differences in means across C and T by separately regressing the binary decision to disclose each secret number over a dummy equal to one for observations belonging to the treatment group. ***, ** and * provide information on the p-value for the estimated coefficient on the treatment dummy, with *** p<0.01, ** p<0.05, * p<0.1. These regressions include subject random effects.

we present the findings of our experiment, first focusing on senders and then on receivers. Participants did not receive any feedback regarding their payoff or the strategy employed by their opponents. Consequently, their behavior in the experiment was solely determined by the beliefs they formed about the strategies of players in the opposite role. As in Jin, Luca and Martin (2021), we consider all the rounds in the analysis.[16] In all regressions, we include subject random effects and, for the analysis of the receivers' data we also cluster standard errors at the session level. Indeed, while senders receive no information during the session, receivers observe the disclosure choices, receiving some indirect information about senders' strategies. However, our qualitative results are robust to alternative specifications.[17]

4.1 Sender's Behavior

Table 1 summarizes disclosure rates across treatments, reporting the average probability that each realization of $\theta \in \Theta$ is disclosed. Constructing this measure is straightforward in control sessions where senders can either disclose or not their secret number. In treatment sessions, where the type is given by two secret numbers, we consider each dimension of $(\theta_1, \theta_2) \in \Theta^2$, with the corresponding disclosure choice of the sender, as a single observation. In this way, we can examine how increasing the size of the information space from one to two secret numbers affects each secret number's disclosure rate.

[16] Restricting attention to the first 30 rounds (before the belief elicitation starts) and/or dropping the first 10 rounds (to account for learning) does not change the main qualitative results of the paper.

[17] In particular, our analysis of the senders' data is robust to clustering standard errors at the session level. We would get equivalent results by removing subject random effects and repeating our analysis clustering standard errors at the subject level.

In control sessions, when a sender's type consists of a single secret number, we find that the probability of disclosing the number is increasing in its value, starting below 6% for the lowest state ($\theta = 1$) and reaching almost one when the secret number is 5—the highest possible type. Contrary to the theoretical prediction, complete unraveling fails to occur in the data. In particular, senders disclose favorable information while withholding unfavorable one.

These patterns are consistent with the findings in Jin, Luca and Martin (2021) providing evidence that our modified design can replicate established results in the literature. Specifically, our receivers can only choose to guess an integer between one and five, whereas Jin, Luca and Martin (2021) allow also $\{1.5, 2.5, 3.5, 4.5\}$ as guesses. Shrinking the action space has the advantage of making it feasible to construct a detailed incentivized belief elicitation task. However, this comes at the cost of allowing for multiple equilibria in the theoretical framework. Thus, the fact that in our control sessions disclosure rates (and also guesses, as we show in Section 4.2) are similar to those found by Jin, Luca and Martin (2021), suggests that our players do not behave according to a potentially different equilibrium.

Table 1 shows that the lowest state ($\theta = 1$) is disclosed 15.7% of the times in the treatment and only 5.9% of the times in the control, but this difference is only statistically significant at the 10% level. Moreover, $\theta = 2$ is disclosed 51.2% of the times in the treatment and only 29.1% of the times in the control and this difference is statistically significant at the 1% level. If we consider the overall disclosure rate of these unfavorable states, we find an increase from 16.5% to 33.4%, which is significant at the 1% level. The disclosure rates for the higher states (3, 4 and 5) do not present any statistically significant difference between treatment and control (and this is true even if we look at the aggregate disclosure rate of these more favorable states). Overall, the probability of disclosing any secret number is higher in the treatment and this difference is significant at the 5% level.

This pattern is confirmed if we compare the distributions of senders' disclosure probabilities in the two treatments. Indeed, Figure 1 shows a first-order stochastic dominance shift of such a distribution for the low secret numbers ($\theta = 1$ and $\theta = 2$) when we increase the dimensionality of the information space. The distributions are, instead, indistinguishable for the higher secret numbers.[18]

Sender - Result 1 (a): Increasing the dimensionality of the information space increases senders' rates of

[18] We perform a two-sample Kolmogorov–Smirnov test. For low secret numbers, the two distributions are significantly different (p-value < 0.01). For high secret numbers, the two distributions are not significantly different from each other.

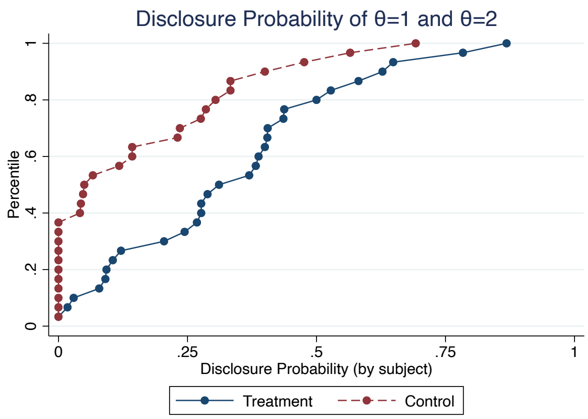

Figure 1: CDFs of Senders' Disclosure Probability

disclosure for unfavorable information.

This analysis only focuses on the disclosure decisions of a given secret number and abstracts from the fact that when presented with a bi-dimensional information space, senders' disclosure decisions can potentially happen at the pair level. In effect, when a sender is endowed with two pieces of information, he can choose whether to reveal nothing, everything, or only one piece, and the same sender may make different disclosure decisions for the same secret number depending on the other secret number drawn.[19] Understanding how the two dimensions jointly affect the disclosure choices is important to disentangle what are the forces that lead to the increase in disclosure in our treatment. In Figure 2, we plot the frequency at which each secret number is disclosed given the value of the other secret number, regardless of whether this other number is disclosed. The reason is that we are only interested in providing evidence that senders do not look at the two dimensions independently rather than documenting the disclosure decision taken for each pair of secret numbers.

Figure 2 shows that the probability of disclosing the lowest state (i.e., 1) or the two highest ones (i.e., 4 and 5) is minimally affected by the value of the others. By contrast, the probability of disclosing a 2 or a 3 is substantially lower when the other secret number is equal to 1 compared to the other possible pairs. In addition, the probability of disclosing a 3 increases when the second secret number is a 4 or a 5.[20] Overall, our results suggest that the probability of disclosing a certain number depends on the value of the other one in the pair, leading senders to disclose numbers that they would not have disclosed with a uni-dimensional information space. This mechanism, which we further explore when analyzing senders' beliefs,

[19]Besides the unraveling equilibrium, our game features the existence of another equilibrium in which the choice of disclosing a 2 or a 3 is affected by the other number. Hence, the theory does not exclude that the disclosure decisions are taken at the pair level.

[20]Most of these differences are statistically significant at a 1% level.

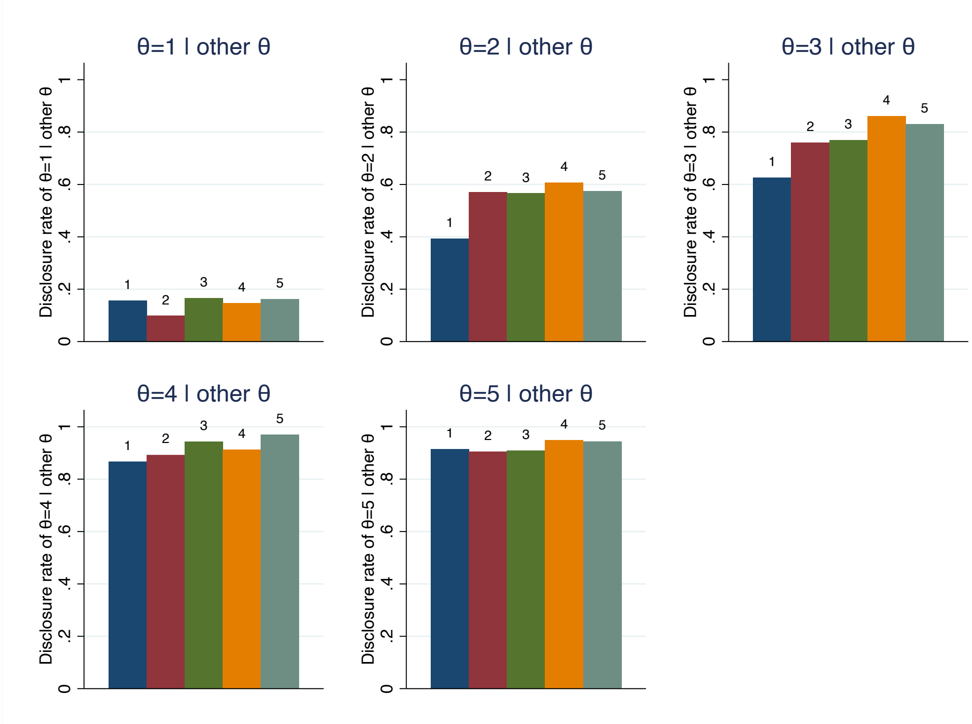

Figure 2: Frequency of Disclosure of a Secret Number Given the Value of the Other - Treatment

Notes: We report the frequencies of disclosure of a given secret number θ for each realization of the other (regardless of its disclosure choice).

is a candidate driver of the increase in the average disclosure rate documented in the treatment relative to the one-number case. Moreover, even if this pattern is potentially consistent with the second type of equilibrium discussed in Section 3, the actual disclosure probabilities deviate significantly from those predicted by such a theoretical equilibrium. This discrepancy indicates that the observed pattern cannot be driven by this equilibrium.[21]

Sender - Result 1 (b): Senders' disclosure choices are taken looking at the pair rather than at the two secret numbers independently. This leads to an increase in the disclosure of low secret numbers.

While our analyses so far shed light on the change in the probability that each type dimension is disclosed and how the disclosure decision is taken after we increase the dimensionality of the information space, a natural policy-related question is what happens to the probability of full disclosure and whether partial disclosure arises in equilibrium. In our Hypothesis 2 about the sender's behavior, we argued that partial disclosure should account for a negligible share of disclosure choices. By contrast, Table 2 shows that partial

[21] In particular, we find significantly less disclosure of the secret number 2 when the other dimension is high and significantly more disclosure of the secret number 3 when the other dimension is low. This implies the presence of substantial deviations from the theoretical predictions of our alternative equilibrium, which we will delve into further in the following part of the paper.

Table 2: Disclosure Choices

	One-number (C)		Two-numbers (T)	
	Frequency	Percent	Frequency	Percent
No Information	642	42.80	248	16.53
Information on A only	·	·	281	18.74
Information on B only	·	·	272	18.13
Full Information	858	57.20	699	46.60

Notes: For each treatment, we report the frequency of each disclosure decision.

disclosure is empirically relevant: in more than 35% of the rounds, senders disclose only one of the two numbers, and in most cases (87.34%), the lowest number of the pair is the one hidden.[22] Also, in roughly half of the instances in which partial disclosure is used, the sender withholds a secret number higher than 1.

Table 2 also shows that when increasing the size of the information space, the probability of full disclosure decreases from 57% to 46.60% and that of no information disclosure from 42.80% to 16.53%. Regressing the probability of partially disclosing one dimension on the difference between such dimension and the undisclosed one, we find that the two variables are positively correlated.[23] Hence, partial disclosure is used by the sender to "hide a flaw" and only disclose the best attribute.

Sender - Result 2: When endowed with two pieces of information, senders engage in partial disclosure.

Finally, we study how senders' disclosure decisions change across rounds and whether there is any evidence of learning. Indeed, even if no feedback is provided, senders might improve their reasoning about the game and adjust their behavior accordingly. Changes in behavior can help us explain changes in disclosure rates across the two treatments: the higher dimensionality of the information space can affect the way in which experience modifies behavior.

Figure 3 reports moving averages of the disclosure rates for each secret number realization, for both control and treatment. As before, for the treatment, each dimension of the bi-dimensional state is considered as a single observation. We can see that the disclosure rates of $\theta = 1$, $\theta = 4$, and $\theta = 5$ are quite stable over time. For $\theta = 1$ the level of the disclosure rate is constantly higher in the treatment, while for the higher states the levels are comparable.[24] For $\theta = 3$, the increase in disclosure rate is steeper in the control

[22] The share 87.34% is computed including the cases in which the two secret numbers are the same and the sender only discloses one of them. The disclosed dimension is strictly higher than the undisclosed one in 77.76% of the times partial disclosure is used.

[23] See the regression reported in Table A.2 in Appendix A.

[24] Overall, the sender's learning patterns in the control are consistent with those in Jin, Luca and Martin (2021) with no feedback

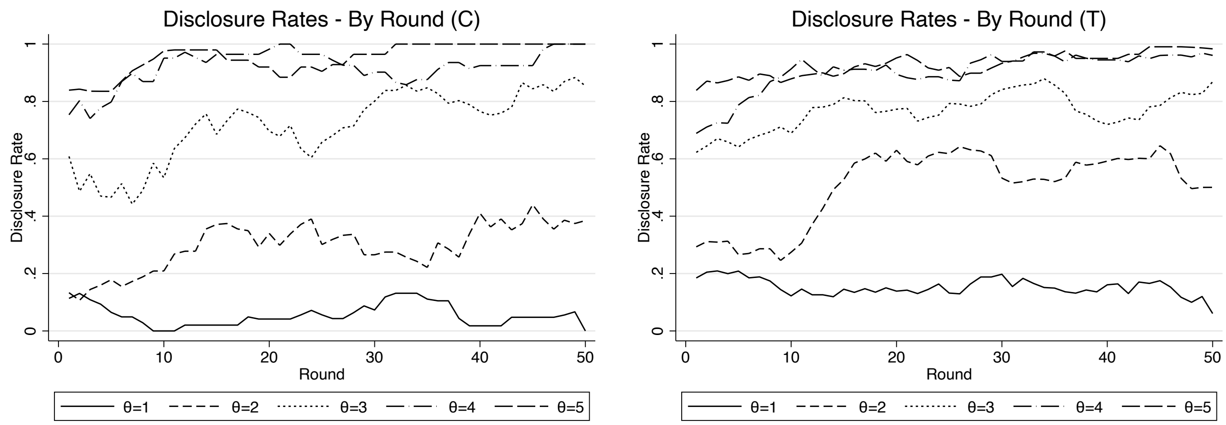

Figure 3: Senders' Disclosure Rates by Round

Notes: For each treatment and each secret number, we plot moving averages of the senders' disclosure rates. Each average is computed using seven rounds: the current round, the previous three rounds, and the subsequent three rounds.

compared to the treatment, but then the disclosure rate stabilizes around the same level. As expected from the previous analysis, the starkest difference lies in the disclosure rates for $\theta = 2$. While the initial levels are similar, after the first 10 rounds we can see a large increase in disclosure rates for the treatment and just a moderate one for the control.[25] Finally, the level at which the disclosure rates stabilize is higher in the treatment, with a difference of roughly 20%. This behavior is consistent with *Result 1 (a)* and *Result 1 (b)*.

4.1.1 Sender's Beliefs

What drives the increase in senders' disclosure rates for low secret numbers? Since we do not provide subjects with any feedback, we conjecture that the shift in senders' behavior is motivated by a change in their beliefs about how receivers will respond to observing that some information is withheld.

First, to verify the validity of our elicited beliefs, we test whether the probability of disclosing a secret number is higher when the average expected guess computed through the sender's beliefs is below its realization. If this is the case, then the sender is effectively best responding to his own belief about the receiver's action. We can verify this by computing the share of times in which senders make a disclosure choice that is optimal given their beliefs about receivers' actions, i.e., they disclose a secret number if it is above what they believe is the guess after no disclosure and they hide it otherwise. This happens roughly 81.5% of the

and fixed roles.

[25]If we look at the distribution of each secret number in each round, we do not find substantial differences that could affect the aggregate disclosure rates over time.

Table 3: Sender's beliefs and disclosure choices

VARIABLES	$\mathbb{1}\{\text{Disclose } \theta\}$
$\mathbb{1}\{\theta >$ Expected guess$\}$	0.151***
	(0.0257)
Round	0.003***
	(0.0004)
Secret Number = 2	0.209***
	(0.0244)
Secret Number = 3	0.488***
	(0.0289)
Secret Number = 4	0.636***
	(0.0306)
Secret Number = 5	0.657***
	(0.0305)
Treatment	0.062**
	(0.0316)
Constant	0.005
	(0.0264)
Observations	4,500
Subjects	60

Notes: *** p<0.01, ** p<0.05, * p<0.1. Each observation is a disclosure decision taken by a sender over a single secret number θ. Expected guess is the sender-specific expected guess given no disclosure. Subject random effects included.

times in the control and 80.7% of the time in the treatment. This already suggests that the elicited beliefs play a role in explaining senders' behavior.[26] However, we also perform an additional—more formal—test by regressing the decision to disclose each secret number on a binary variable which equals one when the secret number is above the sender's average expected guess given no disclosure. In this regression, we also include as covariates the round number, the treatment dummy, secret number fixed effects, and subject random effects (we treat each secret number in the main treatment as an independent observation). Table 3 shows that senders are 15% more likely to disclose when their secret number is above what they expect the receiver will guess if they did not disclose. Moreover, in line with *Sender - Result 1 (a)* and *Sender - Result 1 (b)*, the disclosure rate is on average 6% higher when two secret numbers are drawn.[27]

[26] We also verify that senders' beliefs over receivers' guesses given the decision to disclose a secret number are accurate on average. This suggests that subjects understand what was required by the belief elicitation task and act consistently with such understanding.

[27] We would get qualitatively equivalent conclusions if we used as a regressor the Expected Guess rather than the indicator function $\mathbb{1}\{\theta >$ Expected Guess$\}$. However, the latter specification employs the variable that is directly relevant to the decision of the sender, i.e., whether the secret number is above or below what they expect to be the receiver's response to concealed information.

Next, we analyze how senders' beliefs about receiver guesses vary across disclosure settings. We compare partial disclosure to full concealment, distinguishing whether the sender's private information includes one or two secret numbers. Figure 4 shows that the CDF of the expected guess given no disclosure is similar across treatment and control.[28] Moreover, the average expected guess is 1.98 when senders observe one secret number (control) and 1.97 when they observe two secret numbers, with the difference being statistically indistinguishable from zero.[29]

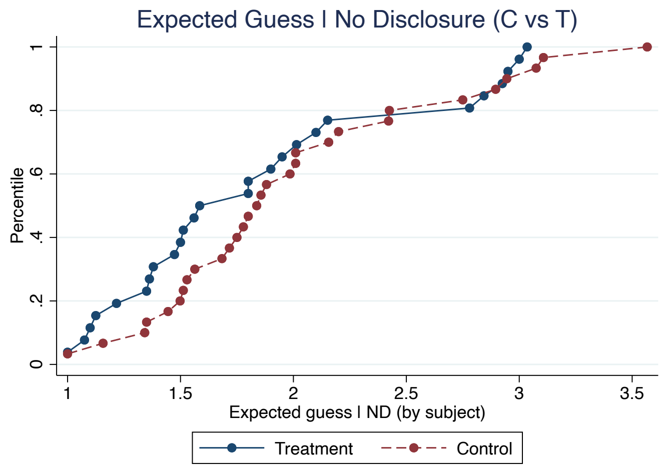

Figure 4: CDF of Senders' Expected Guess given No Disclosure across Treatment and Control

By contrast, senders are more pessimistic about receivers' guesses after observing partial disclosure. The average expected guess drops to 1.75, which is significantly lower (at the 1% confidence level) than that after no disclosure.[30] Focusing on instances in which senders observe two secret numbers, we illustrate in Figure 5 the CDF of the difference in each sender's expected guess if full concealment and partial disclosure are chosen. Figure 5 shows that most senders expect receivers to guess higher numbers if nothing is disclosed, relative to cases in which one secret number is revealed. The average difference is 0.1 and it is significantly different from zero at the 5% level. Moreover, in Figure A.2 in the appendix, we plot the average probability that senders attribute to each receiver's guess after one particular dimension is hidden in the

[28] We perform a two-sample Kolmogorov–Smirnov test and find that the two distributions are not significantly different from each other.

[29] In the appendix, Figure A.1 plots the average probability that senders attribute to each receiver's guess after no information is disclosed, both in the treatment and in the control sessions.

[30] The average expected guess after partial disclosure is 1.93 when the disclosed number is 1, 1.73 when the disclosed number is 2, 1.81 when the disclosed number is 3, 1.74 when the disclosed number is 4 and 1.68 when the disclosed number is 5. In Part II of the experiment, we only have 200 observations wherein partial disclosure arises. Of these, only 19 observations display partial disclosure of secret number 1, and in only 23 we can observe partial disclosure of secret number 2. This suggests that looking at the expected guess given the partial disclosure of each secret number can lead to noisy results. Despite this, our findings demonstrate a clear pattern of increased senders' pessimism regarding guesses after partial disclosure and after no disclosure.

treatment game. Figure A.2 highlights how senders attach up to 60% weight to the probability of guessing 1 when the secret number disclosed is 2.

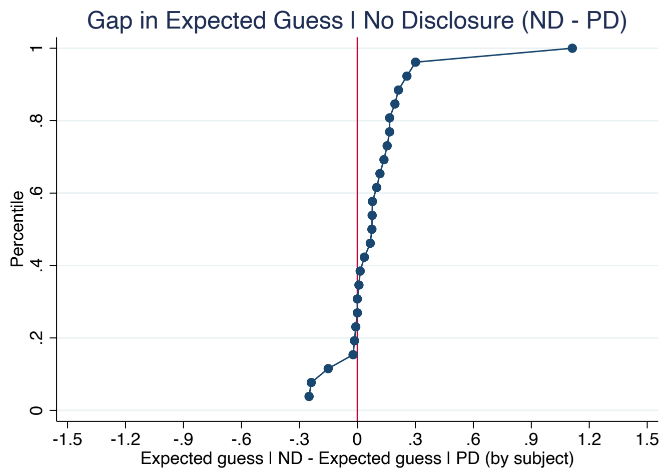

Figure 5: CDF of the Gap in the Expected Guess given No Disclosure and Partial Disclosure

Notes: The figure illustrates the difference in each sender's expected guess given no disclosure and partial disclosure. The sample includes senders in the treatment sessions who have reported their beliefs for at least one instance of no disclosure and partial disclosure.

Overall, our analyses suggest that senders believe that receivers will develop the kind of contingent thinking we conjectured the second dimension would have introduced: *if one dimension is disclosed and the other is not, it must be that the hidden one is low*. In theory, similar reasoning could apply to cases wherein no information is disclosed, but the data show that this is not the case. This suggests that, according to senders, partial disclosure may increase receivers' awareness that disclosure choices are strategic. This possible trait of the receiver's decision when the type-space has two dimensions emerges in our lab experiment because the comparison between different disclosure decisions made by the same player is immediately available to subjects. After partial disclosure, receivers observe a table in which only one number is disclosed, while they see "No Message" in place of the other, and they have to make a guess.[31] Instead, when no information is disclosed, such a comparison is not immediately available, and the receiver needs to compare the current information with the one from previous rounds. This is a much more challenging thought exercise for subjects and may require additional levels of hypothetical thinking that they tend to be unable to perform in these settings (Jin, Luca and Martin, 2021).

It is important to emphasize that the effect on senders' beliefs is the main mechanism that sets our results apart from those in Sheth (2021). In their competitive disclosure game, the presence of a second sender

[31] The interface used in the experiment can be found in Appendix C.

significantly reduces the concealment of evidence compared to the single-sender setting. While senders in both frameworks understand that receivers may compare different disclosure decisions, the key distinction lies in the fact that, in Sheth (2021), these decisions are made by two separate players. As a result, when only one secret number is revealed, receivers can draw almost no inference about a specific sender's strategy. Consistent with this logic, Sheth (2021) finds no significant change in senders' beliefs between the baseline treatment (which matches our control) and the competition treatment: senders do not expect that receivers will become more skeptical of concealed information simply because there are more senders. Instead, they believe receivers will eventually interact more often with senders who disclose private information. This comparison reinforces our intuition that partial disclosure plays a central role in shaping senders' more pessimistic beliefs and, in turn, their strategic behavior in the game.

To sum up, when making their disclosure choice, senders take into account the beliefs that it will induce in receivers. Since they anticipate that partial disclosure—but not no disclosure—raises the receiver's skepticism about undisclosed secret numbers, an increase in the dimensionality of the information space reduces senders' incentives to withhold information. Overall, our findings are consistent with the growth in disclosure rates documented in *Sender - Result 1* being driven by senders' perception of an increased ability of receivers to engage in contingent thinking after partial disclosure messages.

Sender - Result 3: Senders believe receivers will become more skeptical about undisclosed information when they see partial disclosure.

4.2 Receiver's Behavior

While senders believe that in a multi-attribute setting, receivers will be more skeptical about undisclosed information, data say otherwise.

Treating the guess over each secret number as a separate observation, in Table 4 we summarize the average guess both when the secret number is disclosed and when it is not. Looking at each dimension of senders' type as a distinct observation in the treatment, we find that in treatment sessions the average guess upon no disclosure increases up to 2.580 compared to its value of 2.246 in control sessions (and this difference is significant at the 5% confidence level). Figure A.3 in the appendix shows that this pattern is confirmed by the distribution of receivers' guesses. In both treatment and control, the average guess is above the average

Table 4: Summary of Receiver's Guesses

	One-number (C)			Two-numbers (T)		
	Obs	Mean	Std. Dev.	Obs	Mean	Std. Dev.
$\theta = 1$	20	1.100	0.447	96	1.104	0.513
$\theta = 2$	82	2.073	0.409	311	2.148	0.493
$\theta = 3$	227	3.000	0.326	480	3.021	0.408
$\theta = 4$	264	3.848	0.537	514	3.994	0.283
$\theta = 5$	265	4.785	0.665	550	4.909	0.463
ND	642	2.246**	0.844	1,049	2.580**	1.080

Notes: We test differences in means across C and T by separately regressing the guesses taken after each disclosure decision over a dummy equal to one for observations belonging to the treatment group. ***, ** and * provide information on the p-value for the estimated coefficient on the treatment dummy, with *** p<0.01, ** p<0.05, * p<0.1. These regressions include subject random effects and have standard errors clustered at the session level.

value of the undisclosed secret number, with a larger gap in the treatment.[32] Guesses remain, on average, correct when the secret number is disclosed.

Next, we consider whether partial disclosure affects guessing behavior. Table 5 shows that guesses are higher in treatment sessions even when senders provide receivers with partial information.[33] This suggests that, on average, receivers do not treat undisclosed information more pessimistically when one dimension is disclosed, in contrast with what is believed by senders.

Receiver - Result 1: Receivers make, on average, higher guesses about undisclosed secret numbers in the treatment compared to the control.

Similarly to what we did for the sender, we study how guesses after a no-disclosure decision change across rounds. Even if no feedback is provided, receivers are exposed to more, possibly different, information across rounds. Indeed, receivers see the disclosure choices of the senders in the round, and from that they can infer something about the strategies they play. This may induce some learning and affect the way they respond to undisclosed information.

Figure 6 reports moving averages of the guesses receivers make for undisclosed secret numbers, for both

[32]In the control game, the average undisclosed secret number is 1.79. In the treatment game, the average undisclosed secret number in case of no disclosure is 1.91, while in case of partial disclosure is 1.83 — both are higher despite the increase in disclosure given the simultaneous choice of disclosing more often both $\theta_i = 1$ and $\theta_i = 2$. In both cases, the receiver's guesses are incorrect, and the mistake is larger in the treatment.

[33]This change is statistically significant at the 1% level when comparing the control with the partial disclosure of secret number A, while it is not statistically significant when comparing the control with the partial disclosure of secret number B. If we pool together the two instances of partial disclosure, we find that the difference between the guess in the control and the one in the treatment is significant at the 5% level. Comparing the cases of no disclosure, we find a difference that is significant at the 5% level.

Table 5: Average Guess and Partial Disclosure

	One-number (C)			Two-numbers (T)		
	Frequency	Percent	Avg Guess|ND	Frequency	Percent	Avg Guess|ND
No Information	642	42.80	2.246	248	16.53	2.558
Information on A only	·	·	·	281	18.74	2.666
Information on B only	·	·	·	272	18.13	2.529
Full Information	858	57.20	·	699	46.60	·

Notes: For each treatment, we report the frequency of each disclosure decision and the resulting average guess made by the receivers.

control and treatment. For the treatment, we do not distinguish between partial disclosure and no disclosure, but we look at the response to any undisclosed information. In line with Table 5, the average guess after undisclosed information is higher in the treatment than in the control and that is true across the whole sample. In contrast with what was found for senders' disclosure choices, the value of the average guess is highly stable across the 50 rounds, suggesting the absence of learning in receivers' guessing behavior.[34]

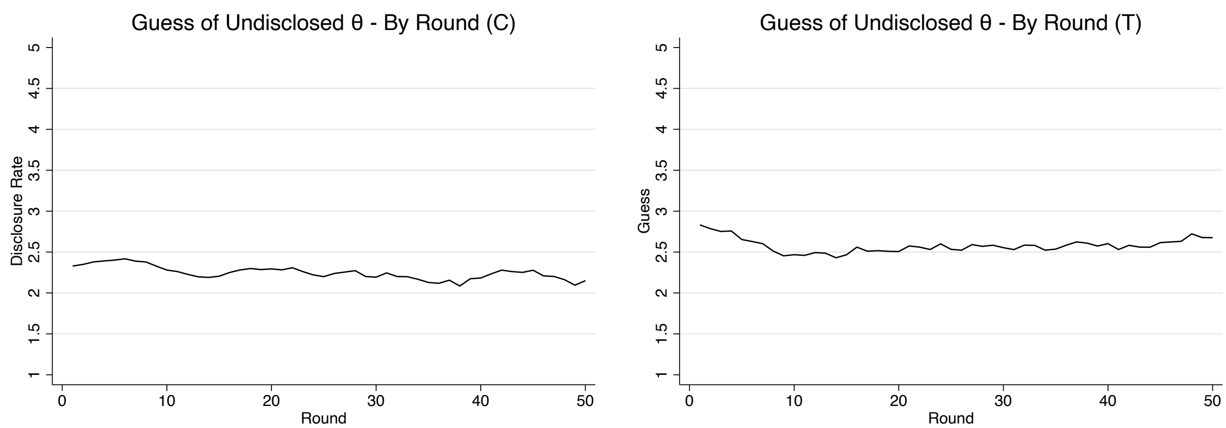

Figure 6: Receiver's Guesses for Undisclosed Secret Number by Round

Notes: For each treatment, we plot moving averages of the receivers' guesses given undisclosed information (both partial. Each average is computed using seven rounds: the current round, the previous three rounds, and the subsequent three rounds.

4.2.1 Receivers' Beliefs

Next, to further shed light on the drivers of receivers' guessing behavior across treatment and control, we analyze their beliefs upon no disclosure.

[34] This pattern is consistent with Jin, Luca and Martin (2021)'s no feedback and fixed roles treatment.

Table 6: Receiver's beliefs and guesses after no disclosure

VARIABLES	(1) Guess	(2) Guess
Expected θ	0.424***	0.425***
	(0.123)	(0.123)
Round	-0.00207	-0.00197
	(0.00202)	(0.00191)
Treatment	0.180**	0.178**
	(0.0699)	(0.0724)
Secret Number = 2		0.0207
		(0.0421)
Secret Number = 3		-0.0443
		(0.0701)
Secret Number = 4		0.0551
		(0.111)
Secret Number = 5		0.0558
		(0.167)
Constant	1.425***	1.417***
	(0.290)	(0.292)
Observations	1,691	1,691
Subjects	60	60

Notes: *** $p<0.01$, ** $p<0.05$, * $p<0.1$. Each observation is a guess taken by a receiver after observing a non-disclosed secret number. Expected θ is the receiver-specific expected secret number given no disclosure. Standard errors clustered at the session level. Subjects' random effects included.

As before, we first verify the validity of our elicited beliefs. This time we do so by investigating whether the receiver's guess increases with the expected secret number computed through the subject's beliefs.[35] Specifically, we leverage the elicited beliefs to compute for each receiver the average expected secret number after observing no disclosure ("Expected θ"). Then, focusing on all instances in which at least one of the secret numbers has not been disclosed, we regress the guess taken by the receiver for that number on that receiver's expected secret number (Expected θ). We control for round number, treatment dummy, subject random effects, and, in column (2), also for secret number fixed effects. Table 6 shows a statistically significant positive correlation between the expected secret number and the guess. Additionally, the regression suggests that on average, receivers tend to guess even more in the treatment relative to the control.[36]

[35]We also study receivers' beliefs over senders' secret numbers given the decision to disclose and we find that those are accurate on average. As before, this suggests that subjects understand what was required by the belief elicitation task and act consistently with such understanding.

[36]With different specifications, the coefficient on the treatment variable is positive, but it may not be significant.

Next, we compute the average expected secret number under no disclosure using the elicited beliefs from both treatments. We find that the average is 2.36 in the treatment group and 1.99 in the control group, with the difference between the two being statistically significant at the 10% level.[37] This suggests that receivers do not become more skeptical of undisclosed information in the treatment relative to the control when no information is disclosed. A more granular analysis of belief distributions under no disclosure reinforces this interpretation: if anything, receivers appear more optimistic in the treatment. As shown in Figure 7, the belief distributions exhibit first-order stochastic dominance, with a statistical significance level of 10%, according to a two-sample Kolmogorov-Smirnov test.

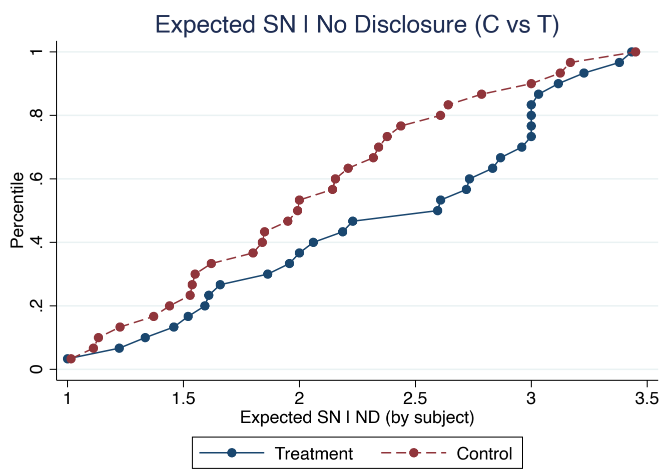

Figure 7: CDF of Receivers' Expected Secret Number given No Disclosure across Treatment and Control

We also analyze the beliefs receivers hold about the non-disclosed secret number given partial disclosure, reported in Figure A.5, and compare them with their beliefs after no information is disclosed. The average expected secret number when receivers see partial disclosure (of any secret number) is 2.44, and is not statistically different from the expected secret number given no disclosure (i.e., 2.36).[38] This indicates that receivers do not evaluate undisclosed information differently in case of partial disclosure and in case of no disclosure. Interestingly, we observe that receivers always place a positive (and relatively large) probability on the event that the undisclosed number is greater than the disclosed one. This is in sharp contrast with the beliefs held by senders and shows that receivers do not internalize the strategic link between the disclosure choices of the two states.

[37] Figure A.4 in the appendix shows the average probability receivers attach to each secret number in these instances.

[38] The average expected secret number after partial disclosure is 2.09 when the disclosed number is 1, 2.33 when the disclosed number is 2, 2.66 when the disclosed number is 3, 2.40 when the disclosed number is 4, and 2.50 when the disclosed number is 5. The same caveat about sample size that we discussed for the analysis of the senders' beliefs applies to this case.

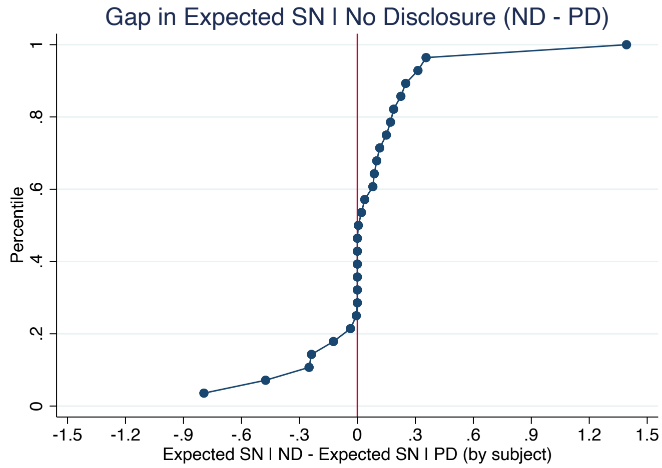

Figure 8: CDF of the Gap in the Expected Secret Number given No Disclosure and Partial Disclosure

Notes: The figure illustrates the difference in each receiver's expected secret number given no disclosure and partial disclosure. The sample includes receivers in the treatment sessions who have reported their beliefs for at least one instance of no disclosure and partial disclosure.

Similar to what we did for the senders, we report in Figure 8 the difference in each receiver's expected secret number after observing full concealment or partial disclosure. A large share of the receivers have non-positive differences, suggesting that they do not assess partial disclosure more pessimistically. This argument is confirmed by the fact that the average difference in the receivers' beliefs is not significantly different from 0.

Receiver - Result 2: Receivers' beliefs about undisclosed information are less skeptical in the treatment compared to the control. This is true for any disclosure strategy of the sender.

This result also has important implications for the sender's optimal disclosure strategy. Senders overestimate the level of sophistication of receivers and best respond to these incorrect beliefs by disclosing more frequently. As a consequence, the disclosure rate for low values of the secret number exceeds the level that would maximize the sender's payoff, given the actual guessing behavior of receivers.

4.2.2 Receivers' Optimal Choice over Two Dimensions

The evidence presented so far suggests that receivers treat the disclosure of the two secret numbers as strategically independent. First, we showed that receivers do not adjust their guesses when senders engage in partial disclosure. Second, we found that increasing the dimensionality of the sender's private information does not make receivers more skeptical of undisclosed attributes. These findings align with Sheth (2021),

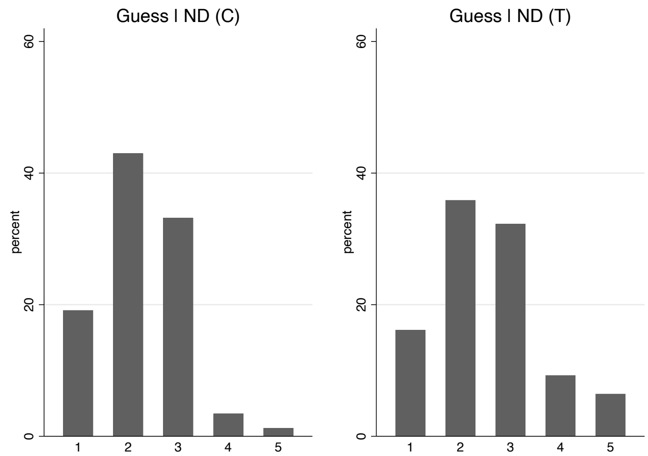

Figure 9: Distribution of Receiver's Guesses Given No Information

Notes: For each treatment, we plot the frequency of receivers' guesses after no information is disclosed. For the main treatment, we combine the guesses made for both secret numbers.

where the two disclosure decisions were made by different competing senders, making it reasonable for receivers to treat them as independent. In our case, however, both disclosure choices are made by the same sender—yet receivers still appear to process them as if they were unrelated. To further examine whether receivers treat the two disclosure choices made by the same sender as independent, we compare the distribution of receivers' guesses across the treatment and the control.

Figure 9 illustrates the distribution of receivers' guesses in the control and the treatment sessions when no information is disclosed. The shape of the distribution remains mostly unchanged, with most of the probability assigned to the values 2 and 3. However, receivers guess 4 and 5 after no disclosure more often in the treatment than in the control. This is consistent with receivers being less pessimistic about undisclosed states in treatment sessions.

Figure 10 plots the distribution of receivers' guesses across the different possible cases in which partial disclosure occurs. When there is partial disclosure of 3, 4, or 5, the shape of the distribution resembles that of Figure 9. In addition, these distributions are very close to the one of the receivers' guesses when no information is provided. By contrast, the distributions resulting from a partial disclosure of 1 or 2 look quite different. A caveat is that because senders often prefer to hide low secret numbers, in our sample, the partial disclosure of secret numbers 1 and 2 only occurs 48 and 72 times, respectively.[39] This generates

[39]The partial disclosure of secret number 3 occurs 122 times, the one of secret number 4 occurs 148 time and that of secret number 5 occurs 163 times. In total, we count 553 observations in which partial disclosure is selected by the sender.

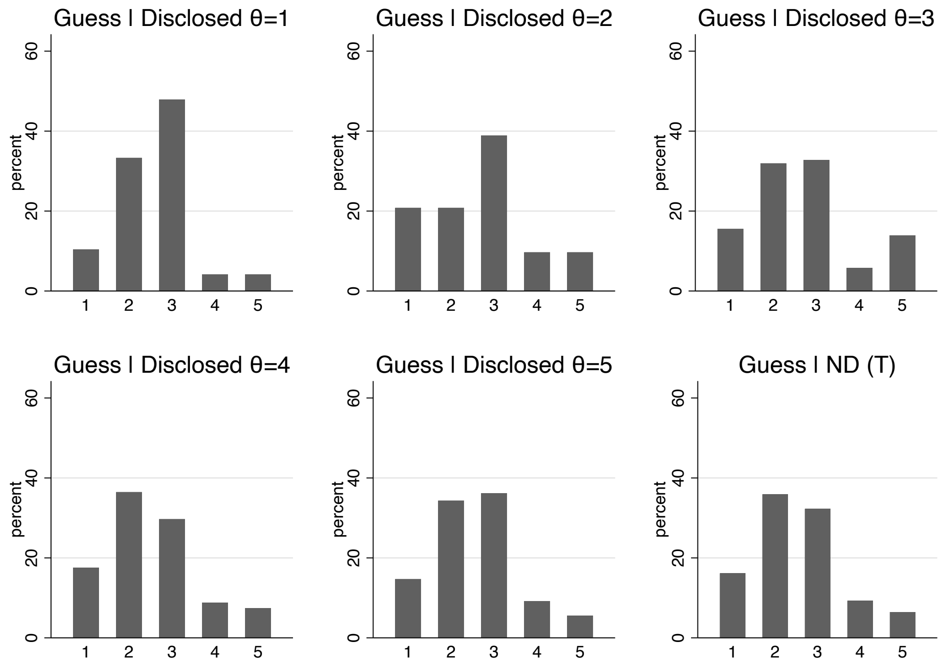

Figure 10: Distribution of Receiver's Guesses Given Partial Information

Notes: We plot the frequency of receivers' guesses to undisclosed information after either partial disclosure or no disclosure in the main treatment. For the partial disclosure cases, we plot the frequency of receivers' guesses conditioning on the disclosed secret number θ. For the no information case, we combine the guesses made for both secret numbers.

a noisy distribution of receivers' guesses. Nonetheless, the main result continues to hold: receivers attach some positive probability to the event that the concealed secret number is greater than the one disclosed.

In net, while we conjectured that observing the partial disclosure of a secret number should trigger the greatest skepticism on the value of the second one, the data indicates that receivers do not internalize that and remain unable to account for the selection of disclosed evidence. Rather, we interpret our empirical evidence as supporting the hypothesis that receivers treat the disclosures of the two secret numbers as strategically independent.

Receiver - Result 3: Receivers' guesses are taken treating the two secret numbers as strategically independent.

5 Conclusion

This paper studies how enabling partial disclosure by increasing the dimensionality of the information space influences receivers' skepticism toward non-disclosed dimensions and, in turn, shapes senders' voluntary disclosure behavior. While Jin, Luca and Martin (2021) document the failure of the unraveling prediction in verifiable disclosure games, we build on a growing literature on contingent reasoning to hypothesize that a richer information environment increases the salience of strategic withholding, thus making receivers more skeptical about concealed evidence and eventually increasing equilibrium disclosure rates.

Through a lab experiment, we show that this conjecture holds in the data. We find that when increasing the number of attributes that a sender can disclose from one to two, the disclosure rate of unfavorable information increases. In particular, the aggregate disclosure rate of 1 and 2 significantly increases from 16.5% to 33.4%, leading to a significant growth in the overall disclosure rate from 57.2% to 65%. By analyzing players' beliefs and receivers' guesses, we shed light on the mechanism driving the increase in disclosure rates, which is the senders' belief that partially disclosing only low secret numbers would facilitate receivers' counterfactual thinking. However, we also demonstrate that these beliefs are, on average, incorrect because, in practice, receivers do not improve in their ability to account for the selection in disclosed dimensions.

Our results contribute to a burgeoning literature on information disclosure by highlighting the role of the dimensionality of the information space and the possibility of partial disclosure in shaping strategic communication. Beyond advancing the understanding of disclosure behavior, the findings point to the importance of considering multi-attribute settings, common in many real-world markets, when interpreting observed disclosure patterns.

References

Agranov, Marina, Jacob K Goeree, Julian Romero, and Leeat Yariv. 2018. "What makes voters turn out: The effects of polls and beliefs." *Journal of the European Economic Association*, 16(3): 825–856.

Aina, Chiara, Andrea Amelio, and Katharina Brütt. 2023. "Contingent belief updating." ECONtribute Discussion Paper.

Ali, S Nageeb, Maximilian Mihm, Lucas Siga, and Chloe Tergiman. 2021. "Adverse and advantageous selection in the laboratory." *American Economic Review*, 111(7): 2152–2178.

Allen, Franklin. 1987. "Discovering personal probabilities when utility functions are unknown." *Management Science*, 33(4): 542–544.

Araujo, Felipe A, Stephanie W Wang, and Alistair J Wilson. 2021. "The times they are A-changing: Experimenting with dynamic adverse selection." *American Economic Journal: Microeconomics*, 13(4): 1–22.

Avoyan, Ala, and Paula Onuchic. 2024. "How Do Groups Speak and How Are They Understood?" working paper.

Barron, Kai, Steffen Huck, and Philippe Jehiel. 2024. "Everyday econometricians: Selection neglect and overoptimism when learning from others." *American Economic Journal: Microeconomics*, 16(3): 162–198.

Benjamin, Daniel J. 2019. "Errors in probabilistic reasoning and judgment biases." *Handbook of Behavioral Economics: Applications and Foundations 1*, 2: 69–186.

Benndorf, Volker, Dorothea Kübler, and Hans-Theo Normann. 2015. "Privacy concerns, voluntary disclosure of information, and unraveling: An experiment." *European Economic Review*, 75: 43–59.

Bertomeu, Jeremy, Paul Ma, and Iván Marinovic. 2020. "How often do managers withhold information?" *The Accounting Review*, 95(4): 73–102.

Beyer, Anne, Daniel A Cohen, Thomas Z Lys, and Beverly R Walther. 2010. "The financial reporting environment: Review of the recent literature." *Journal of accounting and economics*, 50(2-3): 296–343.

Bourveau, Thomas, Matthias Breuer, and Robert Stoumbos. 2020. "Learning to disclose: Disclosure dynamics in the 1890s streetcar industry." *Available at SSRN 3757679*.

Brown, Alexander L, and Daniel E Fragiadakis. 2019. "Benign vs. Self-Serving Information Reduction: Do Individuals Understand the Difference?" *Self-Serving Information Reduction: Do Individuals Understand the Difference*.

Brown, Alexander L, Colin F Camerer, and Dan Lovallo. 2012. “To review or not to review? Limited strategic thinking at the movie box office.” *American Economic Journal: Microeconomics*, 4(2): 1–26.

Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong. 2004. “A cognitive hierarchy model of games.” *The quarterly journal of economics*, 119(3): 861–898.

Charness, Gary, and Dan Levin. 2009. “The origin of the winner’s curse: a laboratory study.” *American Economic Journal: Microeconomics*, 1(1): 207–236.

Chen, Daniel L, Martin Schonger, and Chris Wickens. 2016. “oTree—An open-source platform for laboratory, online, and field experiments.” *Journal of Behavioral and Experimental Finance*, 9: 88–97.

Cohen, Shani, and Shengwu Li. 2022. “Sequential cursed equilibrium.” *arXiv preprint arXiv:2212.06025*.

Danz, David, Lise Vesterlund, and Alistair J Wilson. 2022. “Belief elicitation and behavioral incentive compatibility.” *American Economic Review*, 112(9): 2851–2883.

de Clippel, Geoffroy, and Kareen Rozen. 2020. “Communication, perception, and strategic obfuscation.” *Working Paper*.

Degan, Arianna, Ming Li, and Huan Xie. 2023. “An experimental investigation of persuasion through selective disclosure of evidence.” *Canadian Journal of Economics/Revue canadienne d’économique*, 56(4): 1490–1516.

Deversi, Marvin, Alessandro Ispano, and Peter Schwardmann. 2021. “Spin doctors: An experiment on vague disclosure.” *European Economic Review*, 139: 103872.

Dickhaut, John, Margaret Ledyard, Arijit Mukherji, and Haresh Sapra. 2003. “Information management and valuation: An experimental investigation.” *Games and Economic Behavior*, 44(1): 26–53.

Dranove, David, and Ginger Zhe Jin. 2010. “Quality disclosure and certification: Theory and practice.” *Journal of economic literature*, 48(4): 935–963.

Dye, Ronald A. 1985. “Disclosure of nonproprietary information.” *Journal of accounting research*, 123–145.

Dye, Ronald A, and Mark Finn. 2007. “Equilibrium voluntary disclosures when firms possess random multi-dimensional private information.” *Essays in Accounting Theory in Honour of Joel S. Demski*, 53–72.

Dziuda, Wioletta. 2011. "Strategic argumentation." *Journal of Economic Theory*, 146(4): 1362–1397.

Enke, Benjamin. 2020. "What you see is all there is." *The Quarterly Journal of Economics*, 135(3): 1363–1398.

Esponda, Ignacio, and Emanuel Vespa. 2014. "Hypothetical thinking and information extraction in the laboratory." *American Economic Journal: Microeconomics*, 6(4): 180–202.

Esponda, Ignacio, and Emanuel Vespa. 2018. "Endogenous sample selection: A laboratory study." *Quantitative Economics*, 9(1): 183–216.

Esponda, Ignacio, and Emanuel Vespa. 2023. "Contingent thinking and the sure-thing principle: Revisiting classic anomalies in the laboratory." *Review of Economic Studies*, rdad102.

Eyster, Erik, and Matthew Rabin. 2005. "Cursed equilibrium." *Econometrica*, 73(5): 1623–1672.

Farina, Agata, Guillaume Fréchette, Alessandro Ispano, Alessandro Lizzeri, and Jacopo Perego. 2024. "The selective disclosure of evidence: An experiment." National Bureau of Economic Research.

Fong, Meng-Jhang, Po-Hsuan Lin, and Thomas R Palfrey. 2025. "Cursed sequential equilibrium." *American Economic Review*, 115(8): 2616–2658.

Forsythe, Robert, R Mark Isaac, and Thomas R Palfrey. 1989. "Theories and tests of "blind bidding" in sealed-bid auctions." *The Rand Journal of Economics*, 214–238.

Forsythe, Robert, Russell Lundholm, and Thomas Rietz. 1999. "Cheap talk, fraud, and adverse selection in financial markets: Some experimental evidence." *The Review of Financial Studies*, 12(3): 481–518.

Fréchette, Guillaume R, Alessandro Lizzeri, and Jacopo Perego. 2022. "Rules and commitment in communication: An experimental analysis." *Econometrica*, 90(5): 2283–2318.

Gabaix, Xavier, and David Laibson. 2006. "Shrouded attributes, consumer myopia, and information suppression in competitive markets." *The Quarterly Journal of Economics*, 121(2): 505–540.

Greiner, Ben. 2015. "Subject pool recruitment procedures: Organizing experiments with ORSEE." *Journal of the Economic Science Association*, 1(1): 114–125.

Grossman, Sanford. 1981. "The informational role of warranties and private disclosure about product quality." *Journal of Law and Economics*, 24(3): 461–83.

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

Guttman, Ilan, Ilan Kremer, and Andrzej Skrzypacz. 2014. "Not only what but also when: A theory of dynamic voluntary disclosure." *American Economic Review*, 104(8): 2400–2420.

Hagenbach, Jeanne, and Eduardo Perez-Richet. 2018. "Communication with evidence in the lab." *Games and Economic Behavior*, 112: 139–165.

Heidhues, Paul, Botond Kőszegi, and Takeshi Murooka. 2016. "Inferior products and profitable deception." *The Review of Economic Studies*, 84(1): 323–356.

Hossain, Tanjim, and Ryo Okui. 2013. "The binarized scoring rule." *Review of Economic Studies*, 80(3): 984–1001.

Jehiel, Philippe. 2005. "Analogy-based expectation equilibrium." *Journal of Economic theory*, 123(2): 81–104.

Jehiel, Philippe, and Frédéric Koessler. 2008. "Revisiting games of incomplete information with analogy-based expectations." *Games and Economic Behavior*, 62(2): 533–557.

Jin, Ginger Zhe, and Phillip Leslie. 2003. "The effect of information on product quality: Evidence from restaurant hygiene grade cards." *The Quarterly Journal of Economics*, 118(2): 409–451.

Jin, Ginger Zhe, Michael Luca, and Daniel Martin. 2021. "Is no news (perceived as) bad news? An experimental investigation of information disclosure." *American Economic Journal: Microeconomics*, 13(2): 141–73.

Jin, Ginger Zhe, Michael Luca, and Daniel Martin. 2022. "Complex disclosure." *Management Science*, 68(5): 3236–3261.

Jovanovic, Boyan. 1982. "Truthful disclosure of information." *The Bell Journal of Economics*, 36–44.

Jung, Woon-Oh, and Young K Kwon. 1988. "Disclosure when the market is unsure of information endowment of managers." *Journal of Accounting research*, 146–153.

King, Ronald R, and David E Wallin. 1991*a*. “Market-induced information disclosures: An experimental markets investigation.” *Contemporary Accounting Research*, 8(1): 170–197.

King, Ronald R, and David E Wallin. 1991*b*. “Voluntary disclosures when seller’s level of information is unknown.” *Journal of Accounting Research*, 29(1): 96–108.

Kirschenheiter, Michael. 1997. “Information quality and correlated signals.” *Journal of Accounting Research*, 35(1): 43–59.

Li, Ying Xue, and Burkhard C Schipper. 2020. “Strategic reasoning in persuasion games: An experiment.” *Games and Economic Behavior*, 121: 329–367.

Li, YX, and Burkhard C Schipper. 2018. “Disclosure under unawareness: An experiment.” *University of California, Davis*.

Martínez-Marquina, Alejandro, Muriel Niederle, and Emanuel Vespa. 2019. “Failures in contingent reasoning: The role of uncertainty.” *American Economic Review*, 109(10): 3437–3474.

Mathios, Alan 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.

McKelvey, Richard D, and Talbot Page. 1990. “Public and private information: An experimental study of information pooling.” *Econometrica: Journal of the Econometric Society*, 1321–1339.

Milgrom, Paul. 2008. “What the seller won’t tell you: Persuasion and disclosure in markets.” *Journal of Economic Perspectives*, 22(2): 115–131.

Milgrom, Paul R. 1981. “Good news and bad news: Representation theorems and applications.” *The Bell Journal of Economics*, 380–391.

Montero, Maria, and Jesal D Sheth. 2021. “Naivety about hidden information: An experimental investigation.” *Journal of Economic Behavior & Organization*, 192: 92–116.

Nagel, Rosemarie. 1995. “Unraveling in guessing games: An experimental study.” *The American economic review*, 85(5): 1313–1326.

Niederle, Muriel, and Emanuel Vespa. 2023. “Cognitive limitations: Failures of contingent thinking.” *Annual Review of Economics*, 15: 307–328.

Pae, Suil. 2005. "Selective disclosures in the presence of uncertainty about information endowment." *Journal of Accounting and Economics*, 39(3): 383–409.

Penczynski, Stefan, Christian Koch, and Sihong Zhang. 2022. "Disclosure of verifiable information under competition: An experimental study." *Available at SSRN 4130449*.

Schlag, Karl H, and Joël J Van der Weele. 2013. "Eliciting probabilities, means, medians, variances and covariances without assuming risk neutrality." *Theoretical Economics Letters*, 3(1): 38–42.

Sheth, Jesal D. 2021. "Disclosure of information under competition: An experimental study." *Games and Economic Behavior*, 129: 158–180.

Stiglitz, Joseph E. 2000. "The contributions of the economics of information to twentieth century economics." *The quarterly journal of economics*, 115(4): 1441–1478.

Verrecchia, Robert E. 1983. "Discretionary disclosure." *Journal of accounting and economics*, 5: 179–194.

Verrecchia, Robert E. 2001. "Essays on disclosure." *Journal of accounting and economics*, 32(1-3): 97–180.

Viscusi, W. 1978. "A note on "lemons" markets with quality certification." *Bell Journal of Economics*, 9(1): 277–279.

Appendix

A Additional Figures and Tables

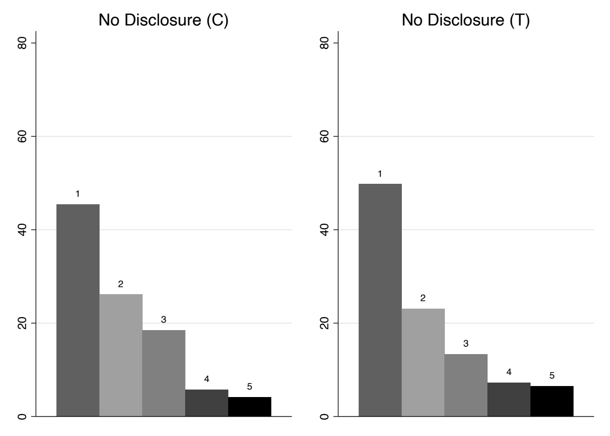

Figure A.1: Sender's Beliefs over Guesses Given No Information - Control and Treatment

Notes: For each treatment, we plot the average probability that senders attached to each receiver's guess after no information is disclosed. For the main treatment, we combine the beliefs elicited for guesses.

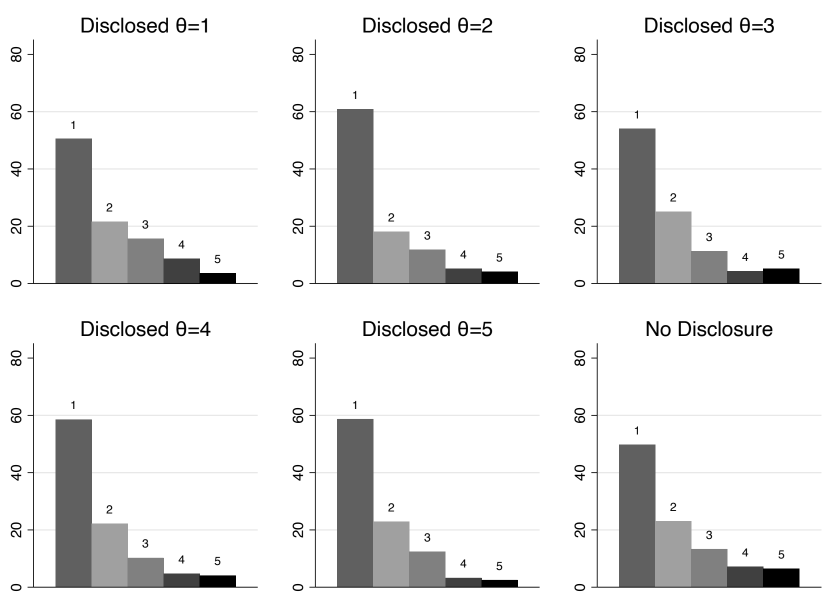

Figure A.2: Sender's Beliefs over Guesses Given Partial Information - Treatment

Notes: We plot the average probability that senders attached to each receiver's guess after either partial disclosure or no disclosure in the main treatment. For the partial disclosure cases, we plot the senders' beliefs conditioning on the disclosed secret number θ. For the no information case, we combine the beliefs elicited for both guesses.

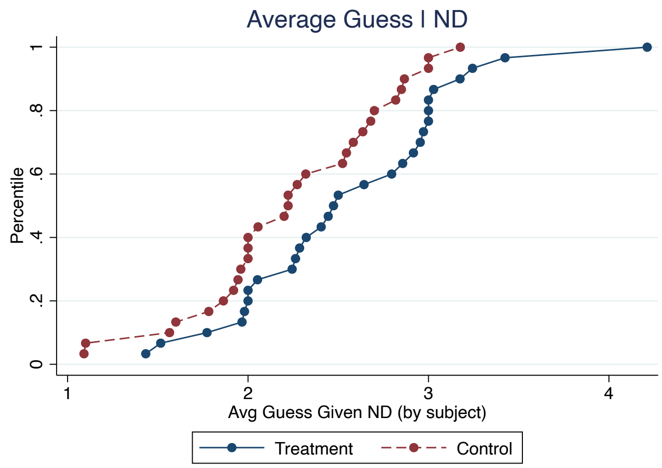

Figure A.3: CDFs of Receivers' Guesses give No Disclosure

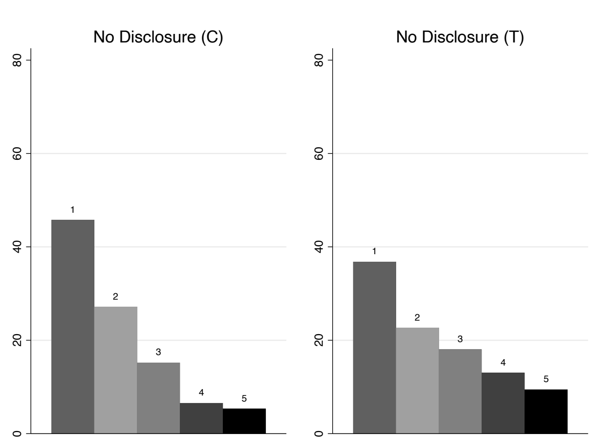

Figure A.4: Receiver's Beliefs over Secret Number Given No Information

Notes: For each treatment, we plot the average probability that receivers attached to each secret number after no information is disclosed. For the main treatment, we combine the beliefs elicited for both secret numbers.

Table A.1: Disclosure Decisions for Each Pair

	No Disclosure		First Disclosed		Second Disclosed		Full Disclosure		Total
	Obs	Percent	Obs	Percent	Obs	Percent	Obs	Percent	Obs
$(1,1)$	54	75.00%	6	8.30%	9	12.50%	3	4.20%	72
$(1,2)$	72	59.50%	6	4.96%	35	28.93%	8	6.61%	121
$(1,3)$	30	25.21%	14	11.77%	64	53.78%	11	9.24%	119
$(1,4)$	7	6.86%	8	7.85%	82	80.39%	5	4.90%	102
$(1,5)$	6	4.84%	5	4.03%	95	76.61%	18	14.52%	124
$(2,2)$	26	39.40%	9	13.64%	3	4.54%	28	42.42%	66
$(2,3)$	25	18.52%	7	5.18%	34	25.19%	69	51.11%	135
$(2,4)$	2	1.96%	12	11.76%	41	40.20%	47	46.08%	102
$(2,5)$	5	4.27%	6	5.13%	47	40.17%	59	50.43%	117
$(3,3)$	7	10.14%	9	13.04%	6	8.70%	47	68.12%	69
$(3,4)$	3	2.40%	5	4.00%	15	12.00%	102	81.60%	125
$(3,5)$	4	3.74%	4	3.74%	17	15.89%	82	76.63%	107
$(4,4)$	3	4.62%	3	4.62%	4	6.15%	55	84.61%	65
$(4,5)$	3	2.78%	3	2.78%	0	0.00 %	102	94.44%	108
$(5,5)$	1	1.47%	2	2.94%	2	2.94%	63	92.65%	68

Notes: In the treatment, for each combination of secret numbers we report the frequency of each disclosure choice. We ignore the order of the secret numbers in the pair.

Table A.2: Correlation between partial disclosure choices and secret number difference

	(1)	(2)
VARIABLES	$\mathbb{1}\{\text{Disclose A only}\}$	$\mathbb{1}\{\text{Disclose B only}\}$
(A-B)	0.0830***	
	(0.005)	
(B-A)		0.0827***
		(0.005)
Observations	1,500	1,500
Subjects	30	30

Notes: *** p<0.01, ** p<0.05, * p<0.1. Subject random effects included.

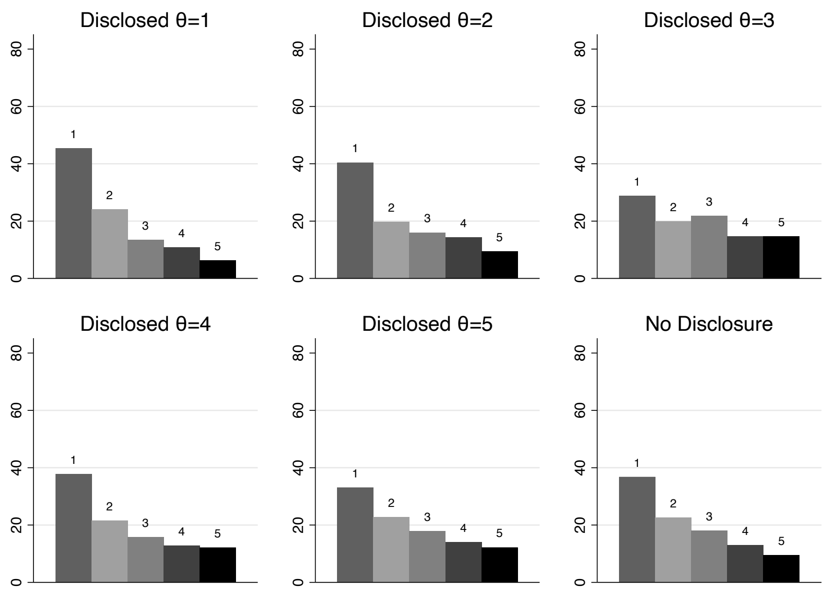

Figure A.5: Receiver's Beliefs over Secret Number Given Partial Information

Notes: We plot the average probability that receivers attached to each sender's secret number after either partial disclosure or no disclosure in the main treatment. For the partial disclosure cases, we plot the receivers' beliefs conditioning on the disclosed secret number θ. For the no information case, we combine the beliefs elicited for both secret numbers.

B Equilibria Multiplicity

B.1 Control Treatment

In the control treatment, the sender is endowed with one state $\theta \in \{1, 2, 3, 4, 5\}$ and he can decide whether to disclose it or not. As mentioned in the main text, there is always a sequential equilibrium in which the sender strictly prefers to disclose every state except the lowest possible one for which he is indifferent. The receiver guesses the right number every time θ is disclosed and guesses $a = 1$ when there is no disclosure. This is clearly a sequential equilibrium and it is consistent with the unraveling equilibrium studied in Milgrom (1981).

However, we assume that $a \in \{1, 2, 3, 4, 5\}$ and this limitation of the action space leads to the existence of another sequential equilibrium. In this equilibrium, the sender strictly prefers to disclose if $\theta \in \{3, 4, 5\}$ and does not disclose if $\theta \in \{1, 2\}$. The receiver guesses the right number when it is disclosed and guesses $a = 2$ when there is no disclosure. Indeed, under our assumption on the payoff and given the uniform distribution over Θ, the receiver is indifferent between guessing $a = 1$ and $a = 2$. When she guesses $a = 2$ with probability 1, the sender strictly prefers to hide $\theta = 1$ and he is indifferent between disclosing and not disclosing $\theta = 2$. This structure clearly represents a sequential equilibrium. However, it is very "fragile". It is enough to add some small probability on the receiver guessing $a = 1$ after no disclosure to make the sender strictly prefer to disclose when $\theta = 2$. Similarly, it is enough to add some small probability to the decision to disclose for the sender if $\theta = 2$ to make $a = 1$ the optimal action for the receiver.

B.2 Main Treatment

In the main treatment, the sender is endowed with a bi-dimensional type $(\theta_1, \theta_2) \in \Theta^2$. There is always a sequential equilibrium in which the sender strictly prefers to disclose any dimension $\theta_i \in \{2, 3, 4, 5\}$ and is indifferent between disclosing any dimension $\theta_i = 1$. The receiver guesses the right number when any dimension is disclosed and $a_i = 1$ every time she sees no disclosure for any component. This equilibrium is the one consistent with the unraveling in Milgrom (1981). However, as in the control treatment, our limited action space allows for additional sequential equilibria.

In another possible sequential equilibrium, the sender discloses his type every time $(\theta_1, \theta_2) \notin \{(1, 1), (1, 2), (2, 1), (2, 2), (1,$

and otherwise fully hides his type. The receiver guesses the right number every time it is disclosed (for any possible dimension). Every time she sees no disclosure, the receiver optimally guesses $a = (2, 2)$. Every time there is partial disclosure, i.e. only one dimension is hidden, the receiver guesses $a_i = 1$ for the hidden θ_i. In this equilibrium, any sender with type $(2, 2)$ is indifferent between fully disclosing and fully hiding his type. In order for such an equilibrium to exist, we need that $\theta = (2, 2)$ decides to disclose his type with a probability smaller than 0.65. This equilibrium predicts less disclosure than the unraveling one and it predicts no partial disclosure (with the exception made for the number 1).

The second possible sequential equilibrium is sustained by an asymmetric behavior of the receiver. In this equilibrium, the receiver guesses $a_j = 2$ when dimension $j \in \{1, 2\}$ is hidden and $a_i = 1$ when dimension $i \neq j$ is withheld. After no disclosure, the receiver guesses $a = (1, 1)$. In response to this receiver's behavior, the sender hides dimension j when $\theta_j \leq 2$, including the case in which $\theta = (2, 2)$. This equilibrium is very fragile. Indeed, if the sender plays a different strategy after $\theta = (2, 2)$, even with an extremely small probability, it cannot be sustained anymore. In addition, the existence of this equilibrium is justified by a strong asymmetric behavior of the receiver, which seems somewhat arbitrary.

Finally, there is an additional sequential equilibrium that follows the idea of the one in the control treatment. The receiver guesses $a_i = 2$ for any dimension that is not disclosed, no matter whether there is partial disclosure or no disclosure. As a response, all senders such that $(\theta_1, \theta_2) \in \{(1, 1), (1, 2), (2, 1), (2, 2)\}$ disclose nothing and all the others hide one dimension i as long as $\theta_i \leq 2$. Notice that, in this setting, the receiver is indifferent between guessing 1 and 2 after no disclosure and the sender is indifferent between disclosing or hiding a dimension equal to 2. As in the control, it is enough to introduce a slight probability that the receiver guesses 1 after no disclosure to change the sender's incentives: any $\theta_i = 2$ would always be disclosed. In the same way, it is enough to add some probability that $\theta_i = 2$ is disclosed by the sender to make the receiver willing to guess $a_i = 1$. As for the control treatment, this kind of equilibrium is extremely fragile to any perturbation of the players' behavior.

B.3 Equilibria Selection

In the previous sections, we analyzed the possible equilibria that can arise in our specific environment. Even if we find that our setting allows for multiple sequential equilibria, we also highlight how the majority of such equilibria rely on the decision of the sender to never disclose if indifferent between the two choices.

While this is a profit-maximizing choice, it can conflict with other types of reasonings. For instance, senders might prefer to be truthful when the choice does not have any direct effect on their payoff: they might have a preference for that or they might be concerned with how the receiver responds to undisclosed information.[1] For this reason, we can more formally select the equilibria of the game by assuming that, when indifferent, the sender discloses their private information with a positive probability ε. We can denote the equilibria surviving this refinement as *"truthful" sequential equilibria*.

It is easy to see that, for the control treatment, the only sequential equilibrium that survives this refinement is the unraveling equilibrium, in which the sender discloses each value of $\theta > 1$ and discloses $\theta = 1$ with at least ε probability. Indeed, the other possible equilibrium, in which $\theta = 2$ is never disclosed, relies on the sender's decision to break the ties towards no disclosure with probability 1, which contradicts our refinement. Hence, the only "truthful" sequential equilibrium is the one in which all information is transmitted.

A similar reasoning applies to the main treatment. It is easy to see that the unraveling equilibrium survives the "truthful" sequential equilibrium refinement since only $\theta_i = 1$ is indifferent between disclosure decisions. The asymmetric equilibria in which the receiver guesses 2 when dimension j is hidden is immediately broken by our refinement. Indeed, the optimality of $a = 2$ after partial disclosure relies on the decision of the sender to always hide $\theta_j = 2$, which is not possible anymore in a "truthful" sequential equilibrium. When $\theta_j = 2$ is disclosed with probability $\varepsilon > 0$, making a guess equal to 2 is not optimal anymore for the receiver, since the scenario in which $\theta_j = 1$ becomes relatively more likely. Similarly, we can argue that the equilibrium in which the receiver guesses $a_i = 2$ after any non-disclosed dimension, no matter whether there is partial or no disclosure, does not survive the refinement. In this equilibrium $(\theta_1, \theta_2) \in \{(1,1),(1,2),(2,1),(2,2)\}$ never discloses any dimension and under any other state realization the dimensions below 2 are hidden as well. Under our refinement $(2,2)$ will disclose both dimensions with a positive probability of at least ε, making $a = (1,1)$ the optimal guess after no disclosure. In the same way, the sender will disclose $\theta_i = 2$ for $i \in \{1,2\}$ with a positive probability, making 1 the optimal guess after partial disclosure. This implies that such equilibrium does not satisfy the "truthful" sequential equilibrium refinement.

We finally need to focus on the equilibrium in which the states

$$(\theta_1, \theta_2) \in \{(1,1),(1,2),(2,1),(2,2),(1,3),(3,1)\}$$

[1] This argument is consistent with the equilibrium refinement used in Fréchette, Lizzeri and Perego (2022).

are fully hidden and the receiver guesses $a = (2,2)$ after seeing no disclosure. The receiver's decision is optimal given the choice of $(2,2)$ to never disclose any dimension, even if indifferent. This equilibrium survives as long as $\varepsilon < 0.65$, so adding a small positive probability to the disclosure of $(2,2)$ is not enough to break the equilibrium.

We can conclude that for the control, there is only one equilibrium outcome surviving the "truthful" refinement, i.e. the full unraveling one. For what concerns the main treatment, two outcomes are serving our refinement: the full unraveling one and the one in which $(\theta_1, \theta_2) \in \{(1,1), (1,2), (2,1), (2,2), (1,3), (3,1)\}$ are fully hidden. In none of these equilibria partial disclosure plays a role.

Online Appendix (For Online Publication Only)

C Instructions and Interface of the Experiment

In this Section, we include the instructions we read to the participants during the experiment and we provide screens of the experimental interface. Figures C.1 and C.5 show the screen used by the sender in the two treatments to make the disclosure choice. Figures C.2 and C.6 show the screen used by the receiver in the two treatments to make a guess. Figures C.3, C.4, C.7 and C.8 show the screens used for the belief elicitation in both treatments.

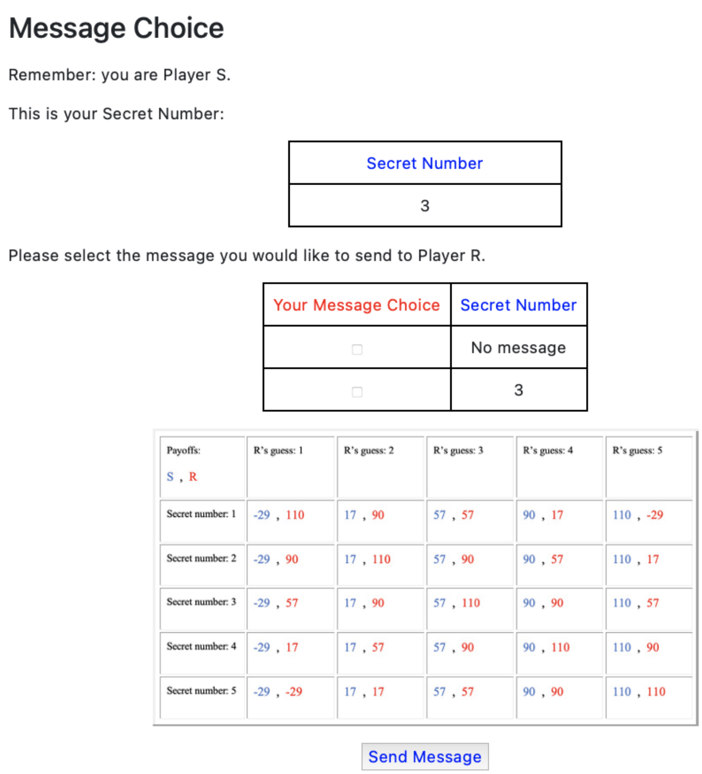

Figure C.1: Sender's Screen in the Control Treatment

Secret Number's Guess

Remember: you are Player R.

Player S sent you the following message:

Secret Number
No message

What is your guess about Player S's Secret Number? [Remember, you can guess 1, 2, 3, 4, 5]

Guess Secret Number

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

Submit Guess

Figure C.2: Receiver's Screen in the Control Treatment

Additional Questions

Someone in this room assigned as Player S was given a Secret Number equal to 4.

This Player S decided to send the following message to their Player R partner:

Secret Number
No message

What is the probability (in %) that this Player R's guess of the Secret Number was equal to each of these values? [Please select a number from 0 to 100 percent.]

Guess	1	2	3	4	5
Probability	%	%	%	%	%

Submit Answer

Figure C.3: Belief Elicitation for the Sender in the Control Treatment

Additional Questions

Someone in this room assigned as Player S decided to send the following message to their Player R partner:

Secret Number
No message

What is the probability (in %) that this Player S's Secret Number was equal to each of these values? Remember, if the above message reports a number, it is this Player S's true Secret Number. [Please select a number from 0 to 100 percent.]

Secret Number	1	2	3	4	5
Probability	%	%	%	%	%

Submit Answer

Figure C.4: Belief Elicitation for the Receiver in the Control Treatment

Message Choice

Remember: you are Player S.

These are your secret numbers:

Secret Number A	Secret Number B
5	5

Please select the message you would like to send to Player R.

Your Message Choice	Secret Number A	Secret Number B
☐	No message	No message
☐	5	No message
☐	No message	5
☐	5	5

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

Send Message

Figure C.5: Sender's Screen in the Main Treatment

Secret Numbers' Guesses

Remember: you are Player R.

Player S sent you the following message:

Secret Number A	Secret Number B
No message	4

What are your guesses about Player S's secret numbers? [Remember, you can guess 1, 2, 3, 4, 5]

Guess Secret Number A	Guess Secret Number B

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

Submit Guesses

Figure C.6: Receiver's Screen in the Main Treatment

Additional Questions

Someone in this room assigned as Player S was given a Secret Number A equal to 2 and a Secret Number B equal to 1.

This Player S decided to send the following message to their Player R partner:

Secret Number A	Secret Number B
2	No message

What is the probability (in %) that this Player R's guess of Secret Number A was equal to each of these values? [Please select a number from 0 to 100 percent.]

Guess of Secret Number A	1	2	3	4	5
Probability	%	%	%	%	%

What is the probability (in %) that this Player R's guess of Secret Number B was equal to each of these values? [Please select a number from 0 to 100 percent.]

Guess of Secret Number B	1	2	3	4	5
Probability	%	%	%	%	%

Submit Answer

Figure C.7: Belief Elicitation for the Sender in the Main Treatment

Additional Questions

Someone in this room assigned as Player S decided to send the following message to their Player R partner:

Secret Number A	Secret Number B
2	No message

What is the probability (in %) that this Player S's Secret Number A was equal to each of these values? Remember, if the above message reports a number for Secret Number A, it is this Player S's true Secret Number A. [Please select a number from 0 to 100 percent.]

Secret Number A	1	2	3	4	5
Probability	%	%	%	%	%

What is the probability (in %) that this Player S's Secret Number B was equal to each of these values? Remember, if the above message reports a number for Secret Number B, it is this Player S's true Secret Number B. [Please select a number from 0 to 100 percent.]

Secret Number B	1	2	3	4	5
Probability	%	%	%	%	%

Submit Answer

Figure C.8: Belief Elicitation for the Receiver in the Main Treatment

Welcome

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

Please silence and put away your cellular phones now.

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

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

Instructions

The experiment you are participating in consists of 50 rounds. Rounds 1 to 30 will be referred as Part I and rounds 31 to 50 as Part II.

At the end of the final round, you will be asked to fill out a questionnaire, which will not affect your earnings in any way. You will then be paid the total amount you have accumulated during the session, in addition to the **$7 participation fee**. Everybody will be paid in private. You are under no obligation to tell others how much you earned.

The currency used during these 50 rounds is what we call Experimental Currency Units (ECU). For your final payment, your earnings during these 50 rounds will be converted into US dollars at a rate of

$$200 \text{ ECU} = \$1.$$

They will then be rounded up to the nearest (non-negative) dollar amount.

Part I

Before the first round, you will be assigned, with equal probabilities, to a **role** (Player S or Player R), and you will play in this role for the entire length of the experiment. In each round, two participants playing in different roles will be **randomly** matched. You are equally likely to be matched with any other person in the room playing in a role different from yours. This means that you will likely play with a different participant in every round. You will not know whom you are matched with, nor will the person who is matched with you.

In each round and for every pair, the computer program will **randomly** generate a Secret Number that is drawn from the set {1, 2, 3, 4, 5}, with each number **equally likely**.

The computer will then show the Secret Number to Player S. After receiving this number, **Player S** will choose whether to report

- No message, or
- The value of the Secret Number

to Player R.

If Player S chooses the "No message" option, Player R will receive **"No message"**.

If the Player S chooses to report the value of the Secret Number, Player R will receive **the actual Secret Number**.

After seeing the message sent by Player S, **Player R** will guess the value of the Secret Number, by inserting a number from the set {1, 2, 3, 4, 5} in the dedicated area. The earnings of both players depend **on the value of the Secret Number and on the Player R's guess**.

The specific earnings (in ECU) are shown in the table below, which is displayed again before Player S and Player R make their choices. Remember:

- in each cell of the table, the payoff for **Player S** is on the **left** (in **blue**), and the payoff for **Player R** is on the **right** (in **red**).
- as you can see from the table, **Player S** earns more when Player R makes a higher guess, and **Player R** earns more when their guess is closer to the Secret Number.

Table: Earnings (in ECU)

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

At the end of each round, your payoff will be computed according to the table above. Your earnings from Part I will be determined by adding up the ECUs you earned in the 30 rounds you will play. You will only know your **total earnings** (from Part I and II) at the end of the session.

Before Part II begins, after completing Round 30, you will receive a message on your screen, after which we will take a break from the experiment and read the instructions for Part II together.

Are there any questions?

We are now ready to start. You will be assigned your role (Player S or R) before the beginning of the game. Once in Round 30, the game will stop, and we will read the instructions for Part II.

Part II

You reached Round 31, where Part II begins.

In Part II, you will keep playing in the **same role** (i.e., either Player S or Player R) as in Part I, and the structure of the game remains the same.

In each round, after the game ends, you will now have to answer a few **additional questions**. How your final earnings from this part are calculated is explained in detail below.

In each round, after your choice as Player S or Player R, you will be shown the message choice of **a Player S selected at random** among the other participants playing in this role:

- If you are **Player R**, you will be asked the probability (expressed as a chance out of 100) that the value of **the Secret Number for this selected Player S** is 1, 2, 3, 4, 5.
- If you are **Player S**, you will be asked the probability (expressed as a chance out of 100) that the **guess of Player R about the Secret Number of this selected Player S** is 1, 2, 3, 4, 5.

This means that, if you are assigned to **Player R** role, you will have to write five probabilities, one for **each possible value of the Secret Number**, and the sum of these five probabilities will have to be 100.

Instead, if you are assigned to **Player S** role, you will have to write five probabilities, one for **each possible value of the guess**, and the sum of these five probabilities will have to be 100.

You will be paid for your choices as in Part I. In addition, the probabilities you have reported in that round will determine your chance of winning a **prize of 100 ECU**. **One round from Part II will be randomly selected for payment.**

Reporting your true guesses about these probabilities is in your best interest **if you want to maximize the earnings from the experiment**. The more accurate your guesses are, the higher the expected payment you are going to receive.

The **exact formula** we will use to pay you is available under request after the experiment. For the sake of brevity, we will not explain it further here. However, it is important to note that if your objective is to maximize the amount of money you are paid in the experiment, then the only way to do that is to enter your true beliefs into the computer when asked.

Are there any questions?

We are now ready to start again. You will keep the same role (Player S or R) assigned to you in Part I. After playing Round 50, the questionnaire will be displayed on your screen. After filling the questionnaire, you will see your total earnings on the screen. Please remain seated at the end of the experiment and wait for instructions from the experimenter.

Welcome

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

Please silence and put away your cellular phones now.

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

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

Instructions

The experiment you are participating in consists of 50 rounds. Rounds 1 to 30 will be referred as Part I and rounds 31 to 50 as Part II.

At the end of the final round, you will be asked to fill out a questionnaire, which will not affect your earnings in any way. You will then be paid the total amount you have accumulated during the session, in addition to the **$7 participation fee**. Everybody will be paid in private. You are under no obligation to tell others how much you earned.

The currency used during these 50 rounds is what we call Experimental Currency Units (ECU). For your final payment, your earnings during these 50 rounds will be converted into US dollars at a rate of

$$200 \text{ ECU} = \$1.$$

They will then be rounded up to the nearest (non-negative) dollar amount.

Part I

Before the first round, you will be assigned, with equal probabilities, to a **role** (Player S or Player R), and you will play in this role for the entire length of the experiment. In each round, two participants playing in different roles will be **randomly** matched. You are equally likely to be matched with any other person in the room playing in a role different from yours. This means that you will likely play with a different participant in every round. You will not know whom you are matched with, nor will the person who is matched with you.

In each round and for every pair, the computer program will **randomly** generate two numbers, "Secret Number A" and "Secret Number B", that are **uniformly** and **independently** drawn from the set {1, 2, 3, 4, 5}. Uniformly means that each number in the set is **equally likely**. **Independently** means that the value of Secret Number A does not affect the value of Secret Number B, and vice-versa (e.g., think about two dice that are rolled at once).

The computer will then show Secret Number A and Secret Number B to Player S. After receiving these numbers, **Player S** will choose whether to report

- No message, or
- The value of Secret Number A only, or
- The value of Secret Number B only, or
- The value of both Secret Number A and Secret Number B.

to Player R.

If Player S chooses not to report any value (meaning, he chooses the No message option), Player R will receive this information: **"No message" for Secret Number A,** and **"No message" for Secret Number B**.

If Player S chooses to report the value of Secret Number A only, Player R will receive this information: the **actual Secret Number A** and **"No message" for Secret Number B**.

If Player S chooses to report the value of Secret Number B only, Player R will receive this information: **"No message" for Secret Number A** and the **actual Secret Number B**.

If the Player S chooses to report the value of both Secret Number A and Secret Number B, Player R will receive this information: the **actual Secret Number A** and the **actual Secret Number B**.

After seeing the information sent by Player S, **Player R** will guess the value of both secret numbers, inserting a number from the set {1, 2, 3, 4, 5} in the space dedicated to the guess of Secret Number A and a number from the set {1, 2, 3, 4, 5} in the space dedicated to the guess of Secret Number B. The earnings of both players depend **on the values of both Secret Numbers and both of Player R's guesses**.

For each pair of secret number and guess, the specific earnings (in ECU) are shown in the table below, which is displayed again before Player S and Player R make their choices. Remember:

- in each cell of the table, the payoff for **Player S** is on the **left** (in **blue**), and the payoff for **Player R** is on the **right** (in **red**).
- as you can see from the table, **Player S** earns more when Player R makes a higher guess, and **Player R** earns more when their guess is closer to the secret number.

Table: Earnings (in ECU) for ***each secret number and guess***

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

In each round, after Player R make their guesses about the two secret numbers, one secret number is **randomly selected**. This secret number and the respective guess will determine your earnings in the round, according to the table above.

Your earnings from Part I will be determined by adding up the ECUs you earned in the 30 rounds you will play. You will only know your **total earnings** (from Part I and II) at the end of the session.

Before Part II begins, after completing Round 30, you will receive a message on your screen, after which we will take a break from the experiment and read the instructions for Part II together.

Are there any questions?

We are now ready to start. You will be assigned your role (Player S or R) before the beginning of the game. Once in round 30, the game will stop, and we will read the instructions for Part II.

Part II

You reached Round 31, where Part II begins.

In Part II, you will keep playing in the **same role** (i.e., either Player S or Player R) as in Part I, and the structure of the game remains the same.

In each round, after the game ends, you will now have to answer a few **additional questions**. How your final earnings from this part are calculated is explained in detail below.

In each round, after your choice as Player S or Player R, you will be shown the message choice of **a Player S selected at random** among the other participants playing in this role:

- If you are **Player R**, you will be asked the probability (expressed as a chance out of 100) that the value of **each secret number for this selected Player S** is 1, 2, 3, 4, 5.
- If you are **Player S**, you will be asked the probability (expressed as a chance out of 100) that the **guess of Player R about each secret number of this selected Player S** is 1, 2, 3, 4, 5.

This means that, if you are assigned to **Player R** role, you will have to write five probabilities **for each secret number** – for a total of ten probabilities – and the sum of these five probabilities will have to be 100.

Instead, if you are assigned to **Player S** role, you will have to write five probabilities **for each Player R's guess** – for a total of ten probabilities – and the sum of these five probabilities will have to be 100.

You will be paid for your choices as in Part I. In addition, the probabilities you have reported in that round will determine your chance of winning a **prize of 100 ECU**. **One round from Part II will be randomly selected for payment.**

Reporting your true guesses about these probabilities is in your best interest **if you want to maximize the earnings from the experiment**. The more accurate your guesses are, the higher the expected payment you are going to receive.

The **exact formula** we will use to pay you is available under request after the experiment. For the sake of brevity, we will not explain it further here. However, it is important to note that if your objective is to maximize the amount of money you are paid in the experiment, then the only way to do that is to enter your true beliefs into the computer when asked.

Are there any questions?

We are now ready to start again. You will keep the same role (Player S or R) assigned to you in Part I. After playing Round 50, the questionnaire will be displayed on your screen. After filling the questionnaire, you will see your total earnings on the screen. Please remain seated at the end of the experiment and wait for instructions from the experimenter.

D Belief Elicitation Details

Given the findings in Danz, Vesterlund and Wilson (2022), we decided to provide subjects with no details about the payment mechanism in the belief elicitation task. However, such details (reported in this Section) were available upon request and, in that case, carefully explained.

D.1 Control

This document shows the exact formula that will be used to determine payments in the additional task of Part II. This document also explains why guessing truthfully maximizes participants' expected earnings. Henceforth, we refer to the guess taken by subjects in the experiment about the decision of the opponent as their "belief."

Let's denote by X the random variable — from the point of view of Player R — describing the realization of the Secret Number given the message choice of the Sender. We will present the reasoning just for the random variable X, but exactly the same reasoning can be extended to the random variable describing the guesses made by Player R given the message choice of Player S.

With the performed task, we have elicited the subject's beliefs about the likelihood of each realization of this random variable. We indicate these beliefs by $\theta = [\theta_1, \theta_2, \theta_3, \theta_4, \theta_5]$, where $\theta_i \in [0, 100]$ for each $i \in \{1, 2, 3, 4, 5\}$.

We want to show that the monetary reward coming from the task in each round is maximized when the subject reports their true beliefs θ. This would in turn imply that reporting beliefs truthfully maximizes the expected final earnings because one round will randomly be selected for payment. The procedure in the experiment can be summarized as follows:

1. The variable X realizes but the subject does not observe such realization;
2. The subject reports the beliefs θ in the computer;
3. The computer draws a number y between 0 and 1 according to a uniform distribution (we denote this random with Y);

4. The loss function
$$l(X,\theta) = \frac{1}{2}\sum_{i=1}^{5}\left(\mathbf{1}_i - \frac{\theta_i}{100}\right)^2$$
is computed, where $\mathbf{1}_i$ is the indicator function equal to 1 if the realization of X is $i \in \{1,2,3,4,5\}$. Notice that $l(X,\theta) \geq 0$ and $l(X,\theta) \leq 1$;

5. The subject wins 100 ECU if $l(X,\theta) < y$ and 0 ECU otherwise.

If the main goal of the subject is to maximize the earnings from this experiment, they will prefer to receive 100 ECU rather than 0 ECU.

We can denote by $P(\theta)$ the probability that the subject receives the prize of 100 ECU. This probability will be such that

$$P(\theta) = \mathbb{E}_Y\mathbb{E}_X[\mathbf{1}_{l(X,\theta)<y}]$$

where the subscript of the expectation indicates the random variable according to which the expectation is computed. If the subject wants to maximize the expected earnings (in ECU), this would mean maximizing

$$\mathbb{E}_Y\mathbb{E}_X[\mathbf{1}_{l(X,\theta)<y}\cdot 100 + \mathbf{1}_{l(X,\theta)\geq y}\cdot 0] = 100\cdot\mathbb{E}_Y\mathbb{E}_X[\mathbf{1}_{l(X,\theta)<y}] + 0\cdot\mathbb{E}_Y\mathbb{E}_X[\mathbf{1}_{l(X,\theta)\geq y}] = P(\theta)\cdot 100.$$

This implies that the final goal of a subject who wants to maximize their earnings from the experiment is to maximize the value of $P(\theta)$, i.e. to select the lottery that assigns the greatest probability to the best prize (i.e., 100 ECU).

Since Y is a uniform distribution with support $[0,1]$ and $l(X,\theta) \leq 1$ for every X and θ, we can conclude that

$$P(\theta) = 1 - \mathbb{E}_X[l(X,\theta)].$$

Hence, to maximize their earning, the subject needs to report the θ that minimizes $\mathbb{E}_X[l(X,\theta)]$. Given the equivalence between these two expressions, we can conclude that

$$\arg\max_{\theta\in\Theta} P(\theta) = \arg\min_{\theta\in\Theta} \mathbb{E}_X[l(X,\theta)]$$

meaning that the θ that minimizes the expected loss function is also the θ associated with the preferred lottery for the agent.

This implies that the constructed mechanism is incentive compatible and telling the true beliefs (i.e. the ones that select the best lottery from the point of view of the subject) is the way to maximize the expected payment from the experiment.

The exact same reasoning can be applied to the elicitation of the beliefs about the guesses of Player R given the message choice of Player S, since now Player S will want to select the best lottery from their point of view.

D.2 Treatment

This document shows the exact formula that will be used to determine payments in the additional task of Part II. This document also explains why guessing truthfully maximizes participants' expected earnings. Henceforth, we refer to the guess taken by subjects in the experiment about the decision of the opponent as their "belief."

Let's denote by X_i the random variable — from the point of view of Player R — describing the realization of the Secret Number $i \in \{A, B\}$ given the message choice of the Sender. We will present the reasoning just for the random variables X_A, X_B, but exactly the same reasoning can be extended to the random variable describing the guesses made by Player R given the message choice of Player S.

With the performed task, we have elicited the subject's beliefs about the likelihood of each realization of these random variables. We indicate these beliefs by $\theta_i = [\theta_{i,1}, \theta_{i,2}, \theta_{i,3}, \theta_{i,4}, \theta_{i,5}]$, where $\theta_{i,k} \in [0, 100]$ for each $k \in \{1, 2, 3, 4, 5\}$ and for each $i \in \{A, B\}$.

We want to show that the monetary reward coming from the task in each round is maximized when the subject reports their true beliefs θ_A, θ_B. This would in turn imply that reporting beliefs truthfully maximizes the expected final earnings because one round will randomly be selected for payment. The procedure in the experiment can be summarized as follows:

1. The variables X_A, X_B realize but the subject does not observe such realization;

2. The subject reports the beliefs θ_A, θ_B in the computer;

3. The computer draws a number y between 0 and 1 according to a uniform distribution (we denote this

random with Y);

4. The loss function

$$l(X_A, X_B, \theta_A, \theta_B) = \frac{1}{4}\sum_{k=1}^{5}\left(\mathbf{1}_{A,k} - \frac{\theta_{A,k}}{100}\right)^2 + \frac{1}{4}\sum_{k=1}^{5}\left(\mathbf{1}_{B,k} - \frac{\theta_{B,k}}{100}\right)^2$$

is computed, where $\mathbf{1}_{i,k}$ is the indicator function equal to 1 if the realization of X_i is $k \in \{1, 2, 3, 4, 5\}$ for $i \in \{A, B\}$. Notice that $l(X_A, X_B, \theta_A, \theta_B) \geq 0$ and $l(X_A, X_B, \theta_A, \theta_B) \leq 1$;

5. The subject wins 100 ECU if $l(X_A, X_B, \theta_A, \theta_B) < y$ and 0 ECU otherwise.

If the main goal of the subject is to maximize the earnings from this experiment, they will prefer to receive 100 ECU rather than 0 ECU.

We can denote by $P(\theta_A, \theta_B)$ the probability that the subject receives the prize of 100 ECU. This probability will be such that

$$P(\theta_A, \theta_B) = \mathbb{E}_Y \mathbb{E}_{X_A, X_B}[\mathbf{1}_{l(X_A, X_B, \theta_A, \theta_B) < y}]$$

where the subscript of the expectation indicates the random variables according to which the expectation is computed. If the subject wants to maximize the expected earnings (in ECU), this would mean maximizing

$$\mathbb{E}_Y \mathbb{E}_{X_A, X_B}[\mathbf{1}_{l(X_A, X_B, \theta_A, \theta_B) < y} \cdot 100 + \mathbf{1}_{l(X_A, X_B, \theta_A, \theta_B) \geq y} \cdot 0] =$$

$$= 100 \cdot \mathbb{E}_Y \mathbb{E}_{X_A, X_B}[\mathbf{1}_{l(X_A, X_B, \theta_A, \theta_B) < y}] + 0 \cdot \mathbb{E}_Y \mathbb{E}_{X_A, X_B}[\mathbf{1}_{l(X_A, X_B, \theta_A, \theta_B) \geq y}] = P(\theta_A, \theta_B) \cdot 100.$$

This implies that the final goal of a subject who wants to maximize their earnings from the experiment is to maximize the value of $P(\theta_A, \theta_B)$, i.e. to select the lottery that assigns the greatest probability to the best prize (i.e., 100 ECU).

Since Y is a uniform distribution with support $[0, 1]$ and $l(X_A, X_B, \theta_A, \theta_B) \leq 1$ for every X and θ, we can conclude that

$$P(\theta_A, \theta_B) = 1 - \mathbb{E}_{X_A, X_B}[l(X_A, X_B, \theta_A, \theta_B)].$$

Hence, to maximize their earning, the subject needs to report θ_A and θ_B that minimizes

$\mathbb{E}_{X_A,X_B}[l(X_A, X_B, \theta_A, \theta_B)]$. Given the equivalence between these two expressions, we can conclude that

$$\arg\max_{\theta_A,\theta_B \in \Theta^2} P(\theta_A, \theta_B) = \arg\min_{\theta_A,\theta_B \in \Theta^2} \mathbb{E}_X[l(X_A, X_B, \theta_A, \theta_B)]$$

meaning that the θ_A, θ_B that minimizes the expected loss function is also the θ_A, θ_B associated to the preferred lottery for the agent.

This implies that the constructed mechanism is incentive compatible and telling the true beliefs (i.e. the ones that select the best lottery from the point of view of the subject) is the way to maximize the expected payment from the experiment.

The exact same reasoning can be applied to elicitation of the beliefs about the guesses of Player R given the message choice of Player S, since now the Player S will want to select the best lottery from their point of view.