Over-monitoring under Transparency Paradox

Natalie Lee*

February 29, 2024

Please click here for the latest version of the paper.

Abstract

Observing others' actions can resolve strategic risk but may also result in second-mover disadvantage and inefficiencies. For example, a high level of transparency in workplaces may enhance operational control but stifle innovation. I theoretically analyze 2-by-2 cyclical games in which a player(manager) selects a probability of monitoring the other(employee)'s action. Employees know the probability but do not know whether their actions are observed. The model predicts the conditions for Transparency Paradox (Bernstein, 2012) in which an intermediate probability of monitoring is optimal for the manager and for the firm's efficiency. If innovation exposes employees to greater potential exploitation by the manager, managers must monitor and exploit with smaller probabilities. However, the laboratory experiment finds that the subjects in the role of manager persistently over-monitor under such a condition, and even those who do not over-monitor tend to over-exploit. In best response, employees refrain from innovation irrespective of the monitoring probability, which incentivizes over-monitoring by managers. Under over-monitoring and over-exploitation, employees bear most of the efficiency loss. The findings highlight the incentive conditions and dynamic processes that make firms and organizations particularly vulnerable to invasive surveillance/monitoring practices that lead to inefficiency and inequality.

1 Introduction

Observing others' actions can provide a substantial strategic advantage in various contexts such as market competition, military conflicts, counter-terrorism, or criminal mon-

*Amsterdam School of Economics, University of Amsterdam. Email: h.lee@uva.nl

itoring. In those contexts, the knowledge of others' actions resolves strategic risk for the observer and facilitates more informed decision-making. The knowledge also enables the observer to influence or control the behavior of the observed: Online platforms collect extensive data on user behavior and profit from their predictive power through targeted ads and recommendation algorithms. At workplaces, new technologies are facilitating ever more pervasive and invasive monitoring and surveillance of employee activities, such as collecting biometric data and remote monitoring and time-tracking (Mateescu & Nguyen, 2019). The digital technology-assisted practices that were adopted during the pandemic have been normalized and stayed with us (Parker et al., 2022; van Dam, 2022). The scope, the intensity, and the depth of surveillance in our markets and organizations raise serious concerns about our rights, freedom, democracy, and equal distribution of power and control (Baiocco et al., 2022; Amnesty International, 2019), yet these concerns are typically overruled by the high economic value and efficiency that can be unlocked by surveillance (Zuboff, 2019).

However, this study will theoretically and experimentally show that unbridled monitoring of others' actions may in fact backfire for the observer as well as economic efficiency. Specifically, observing others' actions resolves strategic risk for the observer, but may also induce second-mover disadvantage and inefficient actions by the observed. For example, Bernstein (2012)'s field study at a mobile-phone manufacturing facility documents *Transparency Paradox* in which high visibility of employee behavior reduces productivity by stifling innovation. Table 1 below shows a 2×2 game between an employee and a manager at a firm as a stylized representation of the problem.

Table 1: Transparency Paradox as a 2×2 game

		Manager	
		Exploit	Status quo
Employee	Innovate	x_1, 190	190 , 100
	Status quo	145 , 10	145 , 100

$x_1 < 145$

An example: Transparency Paradox In this game, the employee decides whether to innovate the production method or produce with the status quo method. The manager

Figure 1: Ex-ante expected payoffs in perfect Bayesian equilibrium

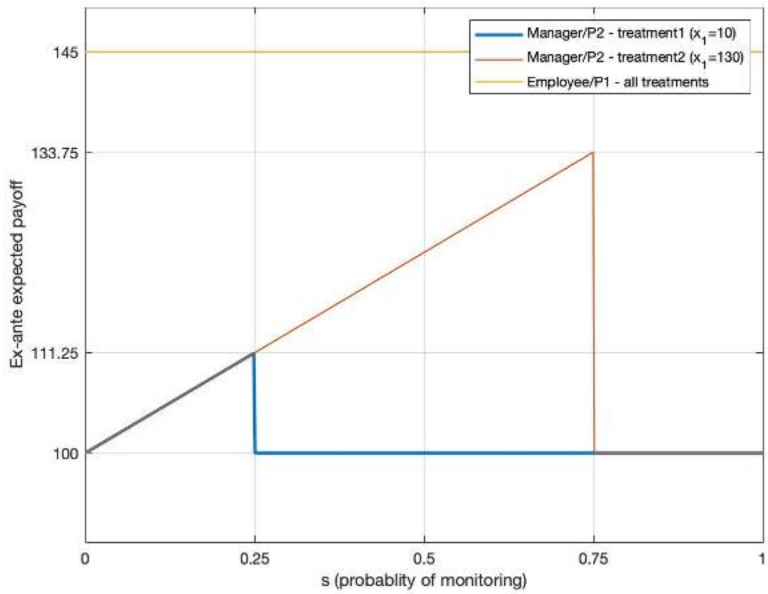

decides whether to exploit[1] the employee or not. Innovation increases efficiency, but the employee is better off innovating only if the manager does not exploit ($x_1 < 145$)[2].

Suppose that the manager chooses a probability of monitoring the employee's action before deciding whether to exploit. The probability of monitoring is common knowledge between the employee and the manager[3], but the employee does not know whether he is indeed monitored or not.

The game presents both incentives and disincentives for monitoring. With a zero probability of monitoring, the cyclical game has a unique Nash equilibrium in non-degenerate mixed strategies, which poses strategic risk that could be resolved through monitoring. Indeed, perfect Bayesian equilibrium (PBE) predicts the manager's ex-ante expected payoff to linearly increase in the monitoring probability, as long as the probability is sufficiently small. Figure 1 plots the expected payoffs for two cases: $x_1 = 10$ and $x_1 = 130$. On the other hand, the manager might not want to monitor with certainty. If she did, the game becomes a sequential-move game in which the employee moves first and the

[1]Exploitation can have multiple interpretations in practice, such as claiming the intellectual property generated by innovation without proper compensation, raising production targets, micromanaging or holding employees accountable for using non-standard production methods (Bernstein, 2012). The example assumes that such actions are likely to be somewhat costly to the manager and unlikely to deliver any benefit unless there is innovation.

[2]The table shows an extreme case in which the employee captures the entire surplus from innovation if the manager does not exploit, but this is not a critical assumption as to be shown in the general model.

[3]Footnote 5 in Section 3 discusses the validity of this assumption.

manager follows. PBE predicts that both players choose status quo and fail to achieve maximal efficiency.

In the unique PBE of this game, the manager maximizes her expected payoff by monitoring with a probability that is strictly between 0 and 1. Specifically, the manager chooses the maximal probability of monitoring that partially resolves strategic risk and also preserves employees' incentives for innovation. The maximal probability of monitoring is 0.25 if $x_1 = 10$, and 0.75 if $x_1 = 130$ as shown in Figure 1.

This study investigates the strategic trade-off between resolution of strategic risk and second-mover disadvantage as demonstrated in the example above. First, I theoretically analyze a class of 2×2 games in which one of the two players (e.g. manager) costlessly monitors the other player (e.g. employee)'s action before taking her own action, with a probability of her choice. The probability is commonly known but the employee does not know whether his action is observed or not. I find the conditions for transparency paradox in which an intermediate probability of monitoring strictly between 0 and 1 is strategically optimal. As in the example, the monitoring player chooses the highest probability that partially resolves the strategic risk without triggering the second-mover disadvantage in the unique perfect Bayesian equilibrium.

Next, a laboratory experiment tests the theoretical prediction that if innovation exposes the employee to greater exploitation by the manager (i.e., smaller x_1 in Table 1), the manager must monitor and exploit with lower probabilities. I find that subjects in the role of manager tend to over-monitor under such condition: They monitor with higher probability compared to the theoretical prediction, and even compared to the alternative condition for which the theory predicted a higher monitoring probability. Furthermore, even those who do not over-monitor tend to over-exploit. In their best response to the over-exploitation, the subjects in the role of employee refrain from innovation irrespective of the monitoring probability. Such an indiscriminate strategy in turn provides strong incentives for over-monitoring by the manager. Over-monitoring and over-exploitation result in a modest but statistically significant efficiency loss, which is mostly borne by the employees. Learning only reinforces over-monitoring. The elicited belief data and play against automated equilibrium strategies suggest that subjects have greater difficulty forming consistent beliefs about the managers' strategy and taking sequentially rational actions as the manager.

One contribution of this study is the simple theoretical model that cleanly captures the strategic trade-off involved in monitoring. Monitoring as a strategic action that creates uncertainty regarding the order of moves has been theoretically explored in relatively few studies, despite its high relevance and broad applicability[4]. Most previous studies consider observability as an exogenous perturbation in games, or in the context of endogenous move games with only second-mover disadvantages but no strategic risk in the theoretical sense. The model in the present paper makes an intuitive prediction that under strategic risk some monitoring is still useful, but the monitoring party should exercise greater discipline if monitoring can be particularly hurtful for the monitored. On the other hand, the experiment show how the monitoring parties fail to show the required discipline, precisely under those circumstances, triggering a chain of strategic responses that strongly incentivizes over-monitoring. The findings suggest that markets and organizations in which the observed party's incentive for efficiency-enhancing action is highly sensitive to over-monitoring may be particularly vulnerable to invasive surveillance practices and resulting inefficiencies and inequality. Finally, this study's findings about managers' monitoring strategy complements Bernstein (2012)'s findings about employees' response to the exogenously varied levels of monitoring.

Although Transparency Paradox at workplaces will continue to serve as the primary example, the class of games featuring both strategic risk and second-mover disadvantage is applicable to other contexts. For example, the two players in Table 1 can be alternatively interpreted as a criminal organization and law enforcement. The law enforcement is the column player who chooses the level of patrol/surveillance. The criminal organization, the row player, chooses between committing a crime and waiting, which respectively correspond to innovating and status quo. The law enforcement chooses between a crackdown and waiting, which respectively correspond to exploiting and status quo.

In what follows, Section 2 reviews relevant previous literature, and Section 3 analyzes a general version of the transparency paradox model. Section 4 describes the experimental design. Section 5 presents findings from the experimental data on behavior and elicited beliefs.

2 Literature

Despite the rapidly growing capabilities of surveillance/monitoring technologies and their broad applicability across contexts, there are relatively few theoretical analyses and even fewer documentations of empirical evidence on monitoring as a strategic choice. To clarify, the following review only covers monitoring of *actions*, which is distinguished from monitoring of exogenous types. In particular, the former creates uncertainty about the order of moves, whereas the latter does not.

[4]Section 2 provides review of the previous literature.

Solan & Yariv (2004) is probably the most closely related study in analyzing the strategic trade-off involved in endogenous monitoring. The authors consider normal-form games in which the second-mover purchases an information device that generates signals about opponent action under a cost structure. Costless monitoring assumed in my model is not precluded by those assumptions, but the choice of monitoring probability for perfectly accurate signal violates their convexity assumption for the information device. Furthermore, in their model, the first-mover chooses his strategy and then second mover makes simultaneous decisions on monitoring and strategy. This assumption implies that the second-mover cannot commit to low-level monitoring, because she cannot lose by purchasing a non-trivial information device if the device is costless. On the other hand, my model examines the second-mover's strategic commitment to low-probability monitoring, and therefore is more relevant to setups where such commitment is credible due to irreversability of the commitment, visibility of monitoring effort (e.g., criminal monitoring through surveillance cameras and partols) or transparency requirement for monitoring policy.

Some previous theoretical studies demonstrated how the exogenous uncertainty in the observability of actions can sharpen equilibrium predictions or offer new interpretations for existing solution concepts (Bagwell, 1995; Van Damme & Hurkens, 1997). For two-person non zero-sum games, Robson (1994) shows existence of Informationally Robust Equilibrium (IRE), a refinement of NE that are robust to exogenous, commonly known probability with which one's pure or mixed strategy, but not the realized action, is revealed to the opponent. Reny & Robson (2004) considers finite two-person games in which the exogenous probability is private knowledge. They show that an equilibrium mixed strategy in the unperturbed game can be approximated by some equilibrium distribution of a perturbed game. Penta & Zuazo-Garin (2022), similar to this study, assumes monitoring that reveals the realized action of mixed strategy. They characterize Rationality and Common Belief in Rationality (RCBR), a refinement of rationalizability that is robust to perturbations in first- or higher-order beliefs about such revelation. Furthermore, they show that the empirical patterns found in asynchronous-move games are consistent with RCBR's equilibrium selection that are advantageous to the first mover.

This study concerns endogenous timing of moves in games, because monitoring can be interpreted as a strategic choice to be a second-mover (follower). Most previous research have studied endogenous timing of moves in the context of market competition where commitment, a strategic choice to be the first-mover(leader), has strategic value. Although theory typically predicts sequential-move Stackelberg outcome rather than simultaneous-move Cournot outcome (Saloner, 1987; Robson, 1990; Hamilton & Slutsky, 1990; Mailath,

Table 2: Game G

		P2	
		A	B
P1	A	x_1, y_4	x_4, y_3
	B	x_3, y_1	x_2, y_2

$x_1 < x_2 < x_4$, $x_1 < x_3$
$y_1 < y_2 < y_4$, $y_3 < y_4$

1993; Ellingsen, 1995; Van Damme & Hurkens, 1999; Amir & Stepanova, 2006), the experiments find more support for the simultaneous-move Cournot outcome (Huck et al., 2002; Fonseca et al., 2006; Müller, 2006). Inequity aversion can explain the Cournot outcome but cannot explain the choice to delay the decision (Santos-Pinto, 2008). Henkel (2002) analyzes a setup in which the first-mover announces its action (price) and simultaneously sets its cost of deviating from the announcement. Under strategic complementarity, the first-mover makes only partial commitment by setting a low cost of deviation, a finding that is qualitatively similar to the intermediate probabilities of monitoring predicted by my model.

3 Model

Consider the 2×2 game in Table 2 in which player 1 (P1) is the row player and player 2 (P2) is the column player. Suppose that before playing the game, P2 chooses a probability of monitoring $s \in [0, 1]$ with which she observes P1's chosen action before P2 takes her own action. With probability s, P2 monitors P1's action, i.e., the game is played as a sequential move game in which P2 is the second mover. If P1 chooses a mixed strategy, P2 monitors the realized action. With the remaining probability $1 - s$, P2 does not monitor, i.e., the game is played in simultaneous moves. P2's choice of s is common knowledge[5], but only P2 knows whether she is informed or uninformed about P1's action. P1 does not

[5]Alternatively, one could assume that s is private knowledge for P2. Under such an assumption, the solution to the current model becomes trivial. There would be no equilibrium in which P1 plays a non-degenerate mixed strategy: if P1 plays a non-degenerate mixed strategy, P2's unique best response is to monitor with certainty ($s = 1$). This is the case if s is assumed to be chosen before all actions but not revealed to P1, or assumed to be chosen simultaneously with all actions. Whereas the private knowledge assumption might better fit some strategic situations, the common knowledge assumption and the reasoning in the present model would be applicable to other situations in which (1) the monitoring parties such as managers and online platforms are required to disclose the extent of their monitoring activities (e.g. by law or to write a contract/ agreement) or (2) the monitoring parties cannot effectively hide its extent of monitoring or (3) the monitoring parties can overrepresent its extent of monitoring (e.g. installing fake speed cameras) but cannot underrepresent it in the long run.

observe P2's action.

I make two assumptions on the payoffs that embody the strategic trade-off in monitoring. First, the game is cyclical: without monitoring, its unique Nash equilibrium in non-degenerate mixed strategy poses strategic risk that can be resolved by monitoring. Without loss of generality, assume $x_1 < x_3$, $x_2 < x_4$, $y_1 < y_2$, and $y_3 < y_4$. Second, monitoring induces second-mover disadvantage: if P2 observes P1's action with certainty ($s = 1$), the resulting sequential-move game has a unique subgame perfect equilibrium which does not yield the first-best outcome for P2. Given the first assumption, P2 matches P1's observed action in any subgame perfect equilibrium. Without loss of generality, assume that P1 plays action B in the equilibrium ($x_2 > x_1$), which does not yield the first-best outcome for P2 ($y_2 < y_4$). All assumptions introduced so far are summarized under Table 2. Now I describe the perfect Bayesian equilibrium of this game and identify the condition for transparency paradox.

Lemma
(a) A subgame following any $s < s^* = \frac{x_4 - x_2}{x_4 - x_1}$ has a unique equilibrium in which

- P1 and non-monitoring P2 plays non-degenerate mixed strategies.
- monitoring P2 matches P1's observed action.
- $\Pi_2(s)$, P2's ex-ante expected equilibrium payoff, linearly increases in s.

(b) A subgame following any $s > s^* = \frac{x_4 - x_2}{x_4 - x_1}$ has a unique equilibrium in which

- P1 and non-monitoring P2 take action B
- monitoring P2 matches P1's observed action
- P1 and P2 respectively get ex-ante expected equilibrium payoffs $\Pi_1(s) = x_2$ and $\Pi_2(s) = y_2$.

(c) The subgame following $s = s^* = \frac{x_4 - x_2}{x_4 - x_1}$ has multiple equilibria in which $\Pi_1(s = s^*) = x_2$. The set of ex-ante expected equilibrium payoffs for P2 includes the expected payoff she would get in case (a) if $s = s^*$, her expected payoff in case (b), and the whole interval between them.

Proposition (Transparency Paradox)

If $s^* = \frac{x_4 - x_2}{x_4 - x_1} > \frac{y_2 - y_3}{y_4 - y_3}$[6],
(a) There exists a unique perfect Bayesian equilibrium in which P2 monitors P1's action with probability $s^* = \frac{x_4 - x_2}{x_4 - x_1} \in (0, 1)$.
(b) In the equilibrium, both P2's ex-ante expected payoff and the sum of P1 and P2's ex-ante expected payoffs are strictly higher than those expected under P2 monitoring with certainty. i.e., $\Pi_2(s^*) > \Pi_2(1)$ and $\Pi_1(s^*) + \Pi_2(s^*) > \Pi_1(1) + \Pi_2(1) = x_2 + y_2$.

The proofs are provided in the appendix. The key example in the introduction provides intuitive explanations for the lemma. The proposition specifies the condition for transparency paradox, i.e., the strategic optimality of monitoring with a probability strictly between 0 and 1.

For an intuitive understanding of the condition, first note that the condition trivially holds for any s^* if $y_2 \leq y_3$. That is, if action A by P1 strictly improves P2's payoff regardless of P2's action, P2 can earn a strictly higher expected payoff than y_2 by monitoring with any probability $s < 1$ as long as it induces P1 to take action A with a strictly positive probability. For example, P2 can simply take action B when non-monitoring and match P1's observed action when monitoring. This strategy strictly improves P2's expected payoff over y_2 regardless of whether she monitors or not. On the other hand, if $y_2 > y_3$, inducing P1 to take action A (reduction in second-mover disadvantage) can backfire if P2 often fails to match the action (increased strategic risk). For P2 to choose $s^* < 1$ in the equilibrium, P1 should be willing to play action A even at a relatively high probability of monitoring $s^* = \frac{x_4 - x_2}{x_4 - x_1}$. This requires small x_2 or large x_1 relative to x_4.

4 Experiment design

The experiment implements 12 rounds of the 2×2 game as shown in Table 3 below, with a pre-game selection of monitoring probability by P2 at the beginning of each round[7].

The experiment has a between-subject design with two treatments. In treatment 1, x_1 set to be 10 such that the theoretically predicted monitoring probability is relatively low: $s^* = \frac{190-145}{190-10} = 0.25$. In treatment 2, x_1 is set to be 130 such that the predicted monitoring probability is relatively high: $s^* = \frac{190-145}{190-130} = 0.75$.

[6] In the opposite case of $s^* = \frac{x_4 - x_2}{x_4 - x_1} \leq \frac{y_2 - y_3}{y_4 - y_3}$, P2's ex-ante expected equilibrium payoff weakly increases in monitoring probability and plateaus at $s \geq s^* = \frac{x_4 - x_2}{x_4 - x_1}$. Therefore, P2 does not face any strict incentive to refrain from monitoring with certainty.

[7] Across the subjects, the two actions for each player are arranged and labeled in counterbalanced order. P1 is always presented as the row player and P2 as the column player to ensure coherence with the instructions.

Table 3: Experimental game

		P2	
		A	B
P1	A	x_1, 190	190 , 100
	B	145 , 10	145 , 100

At the beginning of each experimental session, 12 subjects are randomly selected with equal probabilities to be assigned the role of P1. The remaining subjects are assigned the role of P2. The assigned roles are fixed throughout the session.

In rounds 1 to 10, all subjects play the game against the population (i.e., all subjects in the same session playing the opposite role, Mean-matching protocol - citation) by using the strategy method. Using a slider, P1 specifies a probability with which he would take action A conditional on each possible probability of monitoring chosen by P2. P2 is first asked to select a probability with which she would like to observe P1's action. The instruction makes it clear to P2 that P1 can choose different probabilities of action A depending on the monitoring probability chosen by P2. P2 can choose a probability $s \in \{0.12i | i = 1, 2, ..., 8\} \cup \{0, 1\}$, where 0.24 and 0.72 are the equilibrium predictions for treatments 1 and 2, respectively[8]. The selected probability of monitoring is implemented against all subjects in the role of P1. Once P2 submits the probability of monitoring, she is then asked to specify the probability with which she would take action A in case she does not observe P1's action. In the opposite case where P2 observes P1's action, the experiment automates her best response - matching P1's action.

In rounds 11 and 12, all subjects play against a computerized player that automates the equilibrium strategy of the player in the opposite role. At the beginning of the experiment, subjects are told that there are two additional rounds of the game after the first ten rounds. However, the subjects are not told any other details about those additional rounds. The subjects receive the instructions for rounds 11 and 12 only after they complete the first 10 rounds. The instructions in Appendix X state that the computerized player maximizes the expected number of experimental points under the assumption that all subjects also maximize their own expected number of points assuming computerized player's maximization.

In round 11, subjects play the same game without being told which strategy the computerized player will play. In round 12, the computer screen shows the equilibrium strat-

[8]The discrete choice set ensures a unique equilibrium prediction and allows for elicitation of beliefs conditional on the same set of monitoring probability across the subjects.

egy to be used by the computerized player before the subjects play the game.

In addition to the strategies, the experiment elicits first- and second-order beliefs in rounds 1, 10, and 11. For the first-order beliefs I ask P2[P1] to submit her best guess of the average probability with which P1's[non-observing P2's] took Action A in that round, conditional on each of the ten probabilities of observation P2 could possibly choose. For the second-order, I ask P2[P1] to submit her best guesses about the 10 guesses that have been submitted by one randomly selected P1[P2].

At the end of each round, all subjects receive feedback on their own decisions, the average decisions of other participants in the opposite role, and the number of points they earned in that round. In addition, P1 receive feedback on. P2 receive feedback. The feedback from all previous rounds is displayed in a history table throughout the experiment.

At the end of the experiment, subjects are paid in cash a participation fee of €15 and additional cash prizes they might earn from the strategies and the beliefs they submit.

The strategies in the game are incentivized by a binary lottery with a €10 cash prize. In each round, the subjects earn the expected number of points averaged across all opponents in the session (including the sole computerized player in rounds 11 and 12). At the end of the experiment, one of the 12 rounds is randomly selected for each subject. All rounds are equally likely to be selected. The subject wins the 10 euros with a chance equal to the number of points earned in the selected round divided by 200.

The beliefs are incentivized by the chance of winning a €5 cash prize. The chance of winning determined by one randomly selected belief, according to the binarized scoring rule(Hossain & Okui, 2013). Based on the Danz et al. (2022)'s findings, the instructions simply state that subjects maximize their chance of winning the by truthfully submitting their most accurate beliefs, and the details of the scoring rule is provided in the appendix to the instructions which the subjects are not required to read.

5 Results

The experiment was programmed on Otree (Chen et al., 2016) and conducted at the Center for Research in Experimental Economics and Political Decision-making at the University of Amsterdam. A total of 177 subjects were recruited from the student subject pool. The three sessions in treatment 1 had 30, 28, and 30 subjects each. The other three sessions in treatment 2 had 30, 25, and 34 subjects each.

5.1 P2's monitoring behavior

P2's over-monitor when the theory predicts a low probability of monitoring. Figure 2 shows the mean probability of monitoring chosen by P2's in each round by treatment. In figure 2 and in all other figures to be presented, the intervals around the mean represent 95% confidence interval. The dashed lines represent the equilibrium predictions for each treatment.

Figure 2: Mean probability of monitoring by P2

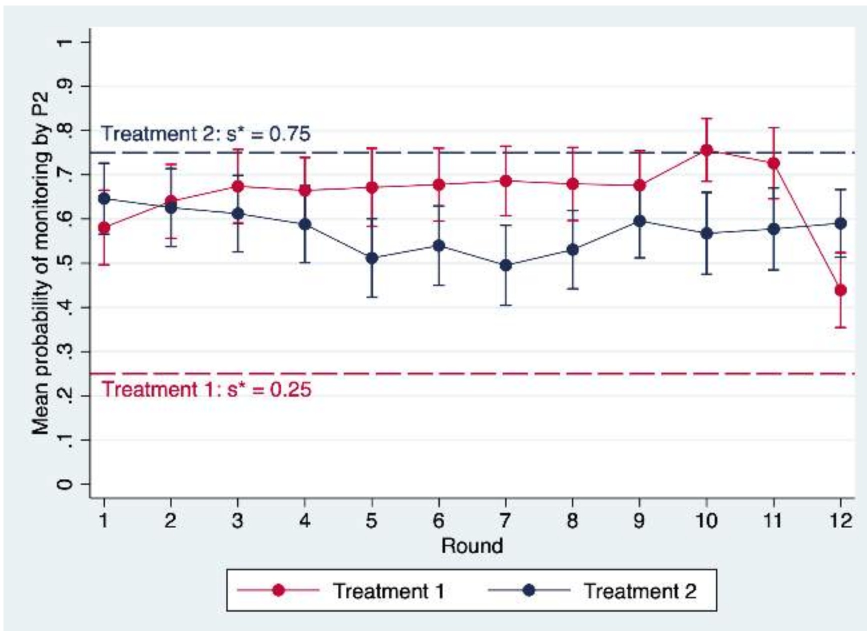

In the first ten rounds in which subjects play against the human population. In these rounds, the mean probability of monitoring in treatment 1 is much higher than 0.25 that was predicted by the equilibrium, and also higher than the observed mean in treatment 2 counter to the theoretical prediction. The mean probability diverges further away from the prediction over the ten rounds (OLS regression of probability of monitoring on round, clustered error at the subject level, $\beta = .0117, p = 0.041$). In treatment 2, P2's probability of monitoring is smaller than predicted. The probability is yet closer to the theoretical prediction compared to treatment 1 and does not exhibit any clear time trend. The difference between the two treatments is statistically significant in round 10, or in all rounds (OLS regression of probability of monitoring on dummy for treatment 1, clustered error at the subject level, $\beta = .0994, p = 0.039$).

In rounds 11 and 12, subjects play against the automated equilibrium strategy. In round 11 in which the automated is unstated, the monitoring probability is not significantly different from that in round 10 in either treatment. In round 12 in which the automated strategy is explicitly stated to subjects, P2's in treatment 1 significantly reduce their monitoring probability compared to round 10 (two-sided ranksum test, $z = -4.002$, $p = 0.0001$). This reduced probability is smaller than the corresponding probability in treatment 2 (two-sided ranksum test, $z = -2.094$, $p = 0.0363$), as predicted by the theory.

The observations in rounds 11 and 12 rules out the explanation that P2's over-monitoring is a rational response to their perceived (lack of) rationality in other subjects. P2's seem to form misspecified beliefs about the equilibrium strategy in treatment 1, and correcting those beliefs brings their behavior closer to the theoretical prediction. Appendix D examines consistency of the elicited beliefs and sequential rationality of the behavior.

5.2 P1's behavior and empirical optimality of over-monitoring

This subsection examines P1's behavior and shows that P2's over-monitoring in treatment 1 is best response to it.

5.2.1 P1's aggregate-level behavior

Figure 3 shows P1's probability of taking action A conditional on P2's probability of monitoring in round 1 and 10 (respectively the first and the last round of play against human players). The step functions in dashed lines represent P1's equilibrium strategy of taking action A only if P2 monitors with a sufficiently low probability, which is supposed to disincentivize P2's monitoring.

P1's indeed decrease their probability of action A in P2's probability of monitoring, as Figure 3 shows[9]. However, P1's mean empirical strategy in treatment 1 is less discriminating than theoretically predicted, and experience only reinforces such a pattern. In round 10, P2's can increase P1's probability of action A by at most 25% points, even if P2 monitors with a zero probability.

5.2.2 Empirical optimality of P2's over-monitoring

The indiscriminating strategy of P1 incentivizes P2 to monitor with a high probability in treatment 1. Specifically, P2's empirical mean payoff increases in her probability of monitoring only in treatment 1. Figure 4 and 5 show the two players' empirically observed mean payoffs in round 8 to 10, conditional on P2's probability of monitoring. P2's mean

[9]OLS regression of P1's probability of action A on P2's probability of monitoring(controlling for round), clustered error at the subject level; $\beta = -.1963$, $p = 0.004$ in treatment 1; $\beta = -.4916$, $p = 0.000$ in treatment 2

Figure 3: Mean probability of action A by P1 conditional on monitoring probability

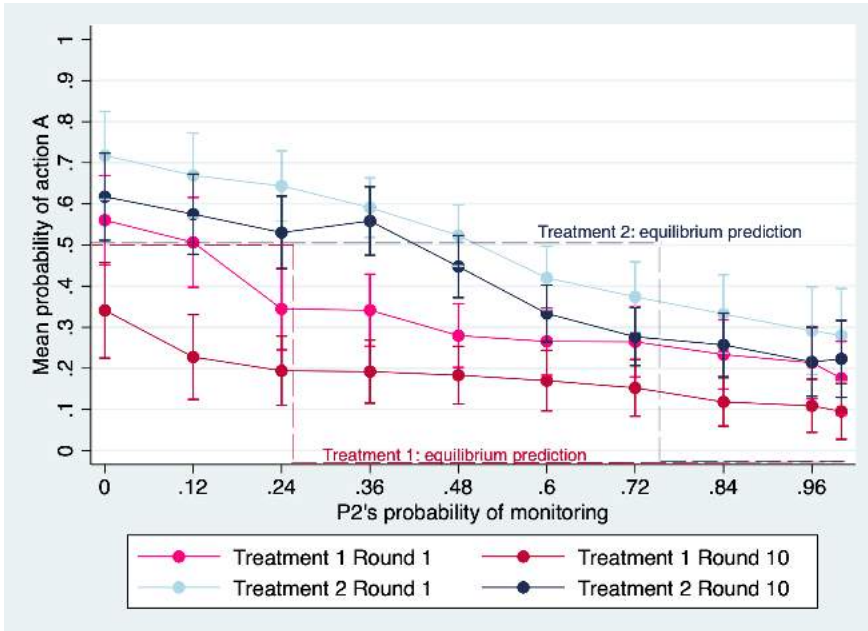

payoffs are represented by the round markers whose values are indicated in the vertical axis on the left. Figure 4 shows a positive relationship between P2's probability of monitoring and her mean payoff in treatment 1, whereas Figure 5 does not show such a relationship in treatment 2[10]. Appendix B applies alternative definitions of mean payoffs to show that qualitatively similar results hold in both early and later rounds.

5.2.3 P1's individual-level behavior

P1's behavior is highly heterogeneous across individual subjects. Figure 9 and 10 in Appendix C provide graphical representation of individual P1's strategy. Although few P1's use a step function-like strategy as theory predicts, some P1's use qualitatively similar strategies. For analytical purpose, define *Discriminating Decreasing Strategy* (DDS) to be a strategy in which the probability of action A weakly decreases in the probability of monitoring, and the average probability of action A conditional on $s < s^*$ is higher than the average probability of action A conditional on $s > s^*$ by a difference of 0.25 or higher. By round 10, 8 out of 36 subjects in treatment 1 and 14 out of 36 subjects in treatment 2 use a DDS. The opposite strategy that is similarly defined (i.e., *Discriminating Increasing Strategy*) is rarely used in either treatment[11].

[10]OLS regression of P2's payoff on probability of monitoring controlling for round, clustered error at the subject level; $\beta = 22.3971$, $p = 0.000$ in treatment 1; $\beta = 2.1392$, $p = 0.401$ in treatment 2

[11]None in treatment 1, round 10; two subjects in treatment 2, round 10

Figure 4: Mean payoff conditional on probability of monitoring, treatment 1

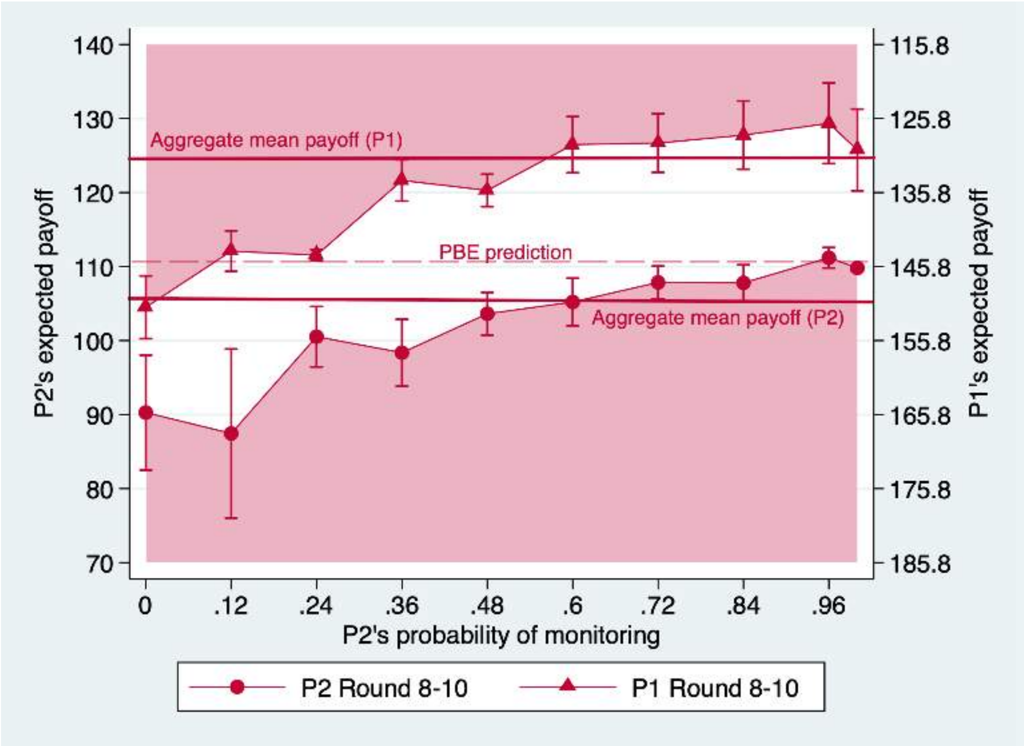

Figure 5: Mean payoff conditional on probability of monitoring, treatment 2

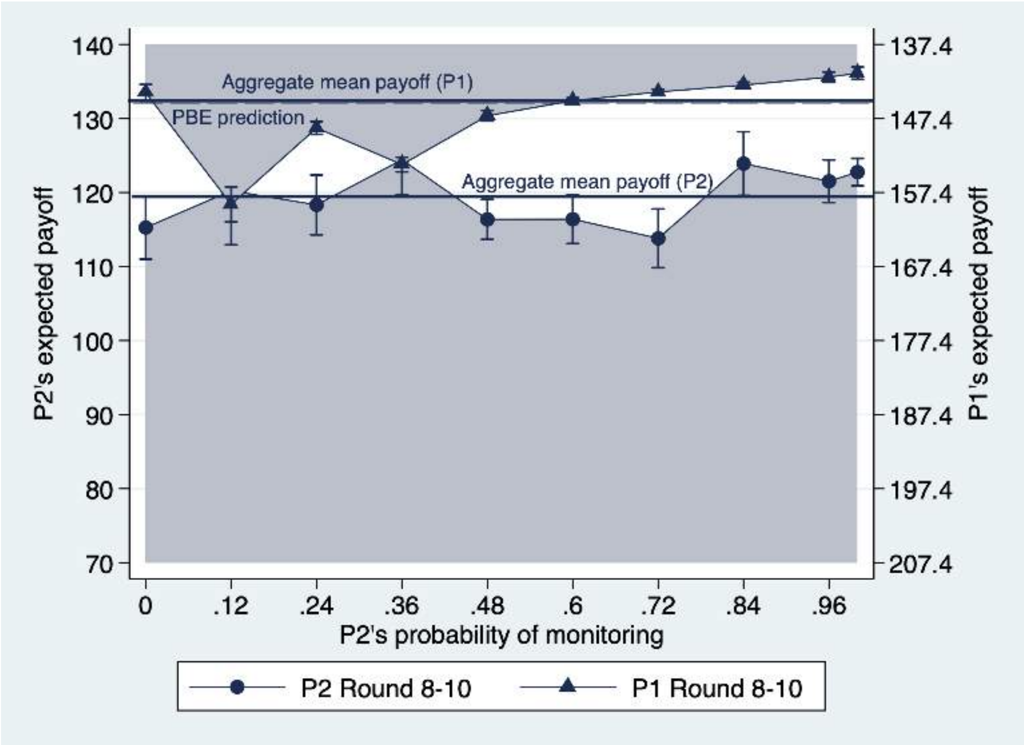

5.2.4 Beliefs about P1's strategy

Both P2's first-order beliefs and P1's second-order beliefs are largely consistent with P1's empirical strategy. Both mean and individual beliefs relate to the monitoring behavior in predicted ways. Appendix D.1 provides graphs and further analyses.

5.3 P2's behavior

Figure 6: Mean probability of action A by P2 conditional on monitoring probability

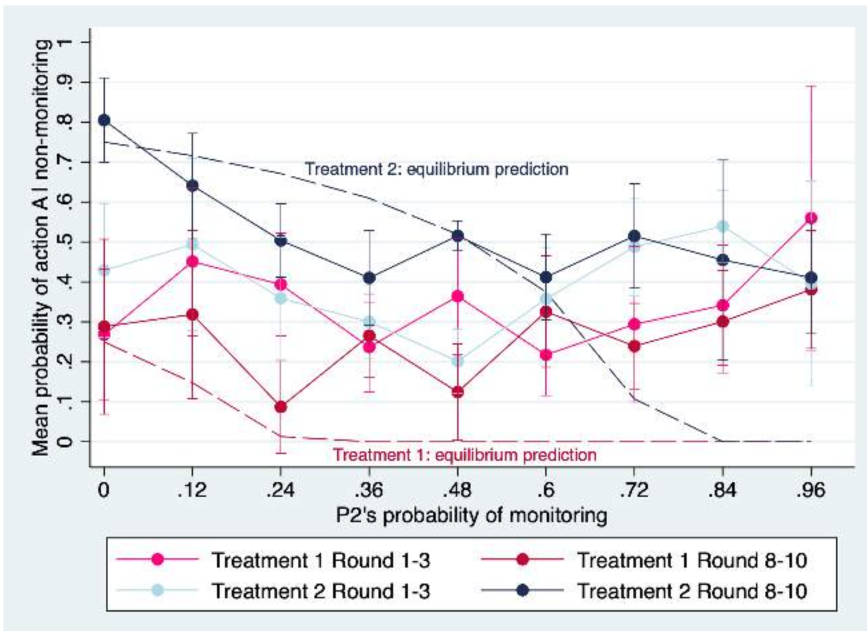

The dashed lines in the figure represent the equilibrium predictions, i.e., the maximum probabilities with which P2 can play action A without triggering P1 to play action B with certainty. Compared to treatment 2, treatment 1 requires a stronger discipline for P2: the non-monitoring type P2's who choose low probabilities of monitoring need to play action A with a much smaller probability in treatment 1 to preserve the incentives for P1 to play action A with a positive probability.

The connected dots in the figure show P2's mean probabilities of action A in the first three and the last three rounds of play against human subjects.

In the first three rounds, the probabilities are similar between the treatments, conditional on low probabilities of monitoring. These probabilities meet the requirement in treatment 1 but not in treatment 2. Even the P2's who choose to monitor with low probabilities in treatment 1 overplay A as the non-monitoring type. This incentivizes P1's to refrain from action A even under a low probability of monitoring, hence the indiscriminate strategy of P1 as shown in figure 3. In treatment 2, the mean strategy by P2 does not count as an overplay of action A.

Figure 6 also shows some indication of learning. By the last three rounds, the mean strategies used by non-monitoring P2's have converged towards the equilibrium predictions in each treatment. The strategies are substantially differentiated between the two treatments conditional on low probabilities of monitoring. In treatment 1, the probability of monitoring conditional on $s < s^* = 0.25$ is not statistically different from the equilibrium prediction.

5.3.1 Beliefs about P2's strategy

Subjects hold onto the inconsistent beliefs about P2's strategy. The mean beliefs are largely constant around 0.5 in both treatments, the individual beliefs are quiet heterogeneous, and there is no sign of learning. The inconsistent beliefs rationalize the non-equilibrium behavior by P1 but do not rationalize over-monitoring or over-exploitation by P2. Appendix D.2 provides graphs and further analyses.

5.4 Efficiency and welfare outcome

Finally, I examine the efficiency and welfare implications. The non-equilibrium behavior observed in the experiment leads to modest efficiency loss in both treatments. In treatment 1 where P2 over-monitors and over-exploits, the efficiency loss is slightly greater and largely borne by the non-monitoring player P1.

First, P1's are less likely to take the efficiency-enhancing action A ("innovate") in treatment 1. As Table 4 shows, P1's take action A with an average probability of 0.16 in treatment 1, and 0.42 in treatment 2, counter to the equilibrium prediction that they would take action A with a half chance in both treatments.

Figures 4 and 5 that were previously introduced plot both players' mean payoffs conditional on P2's probability of monitoring in each treatment. The round markers represent P2's payoffs whose values are indicated on the vertical axis on the left. The triangle markers represent P1's payoffs whose values are indicated on the inverted vertical axis on the right. The perfect Bayesian equilibrium predicts that the two payoffs meet on the dashed line at $s = s^*$. The thick solid lines represent the two player's actual mean payoff aggregated over all probabilities of monitoring. Table 4 summarizes the mean payoffs and comparisons between the two treatments.

Table 4: Efficiency outcome: round 8-10

Role	Outcome	Treatment 1	Treatment 2	Ranksum test comparison
P1	probability of action A [prediction]	0.16 [0.50]	0.42 [0.50]	p=0.000
	mean [prediction]	131.6 [145]	144.0 [145]	p=0.000
	(share in %)	(91%)	(99%)	p=0.000
P2	mean [prediction]	105.1 [110.8]	119.5 [132.4]	p=0.000
	(share in %)	(95%)	(90%)	p=0.070
Total	mean [prediction]	236.7 [255.8]	263.5 [277.4]	p=0.000
	(share in %)	(93%)	(95%)	p=0.009

The figures and the table show efficiency loss that amounts to 7% of the total equilibrium payoffs (i.e., sum of two players' payoffs) in treatment 1, and 5% in treatment 2. Whereas the efficiency loss in treatment 2 is almost entirely driven by P2's loss relative to the equilibrium prediction, the loss in treatment 1 is largely driven by P1 who on average suffer a 10% loss. Note that this loss is entirely borne by the small number of P1's who take the efficiency-enhancing action A.

6 Conclusion

This study theoretically and experimentally investigated the strategic trade-off involved in monitoring other's actions. The theoretical analysis identify the conditions for transparency paradox in which monitoring with an intermediate probability is strategically optimal. The laboratory experiment finds the evidence of such strategic suppression of monitoring. However, when the strategic situation requires a lower probability of monitoring to preserve the incentives for efficient actions by the other, the monitoring player fails to show sufficient discipline. The population consequently converges to the inefficient and un non-equilibrium outcome in which the monitoring player over-monitors and the other player refrains from the efficient action.

The findings suggest that certain markets and organizations might be more susceptible to invasive surveillance practices and resulting inefficiencies. Specifically, institutions that rely on the efficiency-enhancing act of those under surveillance, and those acts are performed at greater risk, might be more likely to end up with inefficient invasive practices.

Future studies could consider alternative strategic situations such as those with bilateral monitoring, costly monitoring, or simultaneous choice of monitoring and other strategic actions.

References

Amir, R., & Stepanova, A. (2006). Second-mover advantage and price leadership in bertrand duopoly. *Games and Economic Behavior*, *55(1)*, 1-20.

Bagwell, K. (1995). Commitment and observability in games. *Games and Economic Behavior*, *8(2)*, 271- 280.

Baiocco, S., Fernández-Macías, E., Rani, U., & Pesole, A. (2022). The algorithmic management of work and its implications in different contexts. *Seville: European Commission*, *JRC129749*.

Barrachina, A., Tauman, Y., & Urbanod, A. (2014). Entry and espionage with noisy signals. *Games and Economic Behavior*, *83*, 127 - 146.

Bernstein, E. S. (2012). The transparency paradox: a role for privacy in organizational learning and operational control. *Administrative Science Quarterly*, *57*, 181-216.

Bernstein, E. S., & Turban, S. (2018). The impact of the 'open' workspace on human collaboration. *Philosophical Transactions B*, *373*, 20170239.

Bruns, B. R. (2015). Names for games: Locating 2×2 games. *Games*, *6*, 495-520.

Chen, D. L., Schonger, M., & Wickens, C. (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, *9*, 88-97.

Danz, D., Lise, V., , & Wilson, A. J. (2022). Belief elicitation and behavioral incentive compatibility. *American Economic Review*, *112 (9)*, 2851-2883.

Ellingsen, T. (1995). On flexibility in oligopoly. *Economics Letters*, *48(1)*, 83-89.

Fonseca, M. A., Müller, W., & Normann, H. T. (2006). Endogenous timing in duopoly: Experimental evidence. *International Journal of Game Theory*, *34(3)*, 443-456.

Hamilton, J. H., & Slutsky, S. M. (1990). Endogenous timing in duopoly games: Stackelberg or cournot equilibria. *Games and Economic Behavior*, *2(1)*, 29-46.

Henkel, J. (2002). The 1.5th mover advantage. *The RAND Journal of Economics*, *33(1)*, 156–170.

Hossain, T., & Okui, R. (2013). The binarized scoring rule. *The Review of Economic Studies*, *80 (3)*, 984 - 1001.

Huck, S., Müller, W., & Normann, H. (2002). To commit or not to commit: Endogenous timing in experimental duopoly markets. *Games and Economic Behavior*, *38(2)*, 240–264.

Laffont, J.-J., & Tirole, J. (1993). *A theory of incentives in procurement and regulation*. MIT Press.

Mailath, G. (1993). Endogenous sequencing of firm decisions. *Journal of Economic Theory*, *59(1)*, 169–182.

Mateescu, A., & Nguyen, A. (2019). Workplace monitorin and surveillance. *Data & Society*. Retrieved from `https://datasociety.net/library/explainer-workplace-monitoring-surveillance/`

Müller, W. (2006). Allowing for two production periods in the cournot duopoly: Experimental evidence. *Journal of Economic Behavior & Organization*, *60(1)*, 100–111.

Parker, K., Horowitz, J. M., & Minkin, R. (2022). Covid-19 pandemic continues to reshape work in america. *Pew Research Center*. Retrieved from `https://www.pewresearch.org/social-trends/2022/02/16/covid-19-pandemic-continues-to-reshape-work-in-america/`

Penta, A., & Zuazo-Garin, P. (2022). Rationalizability, observability and common knowledge. *Review of Economic Studies*, *89*, 948-975.

Rapoport, A., & Guyer, M. (1966). A taxonomy of 2×2 games. *General Systems*, *1*, 203-214.

Reny, P. J., & Robson, A. J. (2004). Reinterpreting mixed strategy equilibria: a unification of the classical and bayesian views. *Games and Economic Behavior*, *48*, 355-384.

Robson, A. J. (1990). Duopoly with endogenous strategic timing: Stackelberg regained. *International Economic Review*, *31(2)*, 263-274.

Robson, A. J. (1994). An informationally robust equilibrium for two-person nonzero-sum games. *Games and Economic Behavior*, *7*, 233-245.

Saloner, G. (1987). Cournot duopoly with two production periods. *Journal of Economic Theory*, *42(1)*, 183-187.

Santos-Pinto, L. (2008). Making sense of the experimental evidence on endogenous timing in duopoly markets. *Journal of Economic Behavior & Organization*, *68(3-4)*, 657–666.

Solan, E., & Yariv, L. (2004). Games with espionage. *Games and Economic Behavior*, *47*, 172-199.

Amnesty International. (2019). *Surveillance giants: How the business model of google and facebook threatens human rights.*

van Dam, A. (2022). The remote work revolution is already reshaping america. *Washington Post*. Retrieved from `https://www.washingtonpost.com/business/2022/08/19/remote-work-hybrid-employment-revolution/`

Van Damme, E., & Hurkens, S. (1996). Commitment robust equilibria and endogenous timing. *Games and Economic Behavior*, *15*, 290-311.

Van Damme, E., & Hurkens, S. (1997). Games with imperfectly observable commitment. *Games and Economic Behavior*, *21*, 282-308.

Van Damme, E., & Hurkens, S. (1999). Endogenous stackelberg leadership. *Games and Economic Behavior*, *28(1)*, 105-129.

Walker, J. M., Smith, V. L., & Cox, J. C. (1990). Inducing risk-neutral preferences: An examination in a controlled market environment. *Journal of Risk and Uncertainty*, *3*, 5-24.

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. New York: PublicAffairs.

Appendix A Proofs

A.1 Proof of Lemma

(a) First, rule out all equilibria in which P1 plays a pure strategy. If P1 takes action A with probability 1, P2 best responds with action A regardless of whether she monitors of not, and P1 gets a payoff of x_1. Then P1 can strictly increase his expected payoff by deviating to action B, regardless of whether P2 monitors or not. If P1 plays B with probability 1, P2 best responds with action B regardless of whether she monitors of not, and P1 gets a payoff of x_2. If P1 deviates to action A, his expected payoff is $sx_1 + (1-s)x_4$, which is strictly greater than x_2 given that $s < \frac{x_4-x_2}{x_4-x_1}$.

Next, consider equilibria in which P1 plays a non-degenerate mixed strategy. Suppose that P1 and non-monitoring P2 play action A with probabilities $p \in (0,1)$ and $q \in [0,1]$ respectively. Recall that monitoring P2's best response is to match the observed action of P1.

P1 plays a non-degenerate mixed strategy if and only if

$$sx_1 + (1-s)(qx_1 + (1-q)x_4) = sx_2 + (1-s)(qx_3 + (1-q)x_2), \tag{1}$$

which implies that non-monitoring P2's probability of action A is

$$q = \frac{x_4 - x_2 - s(x_4 - x_1)}{(1-s)(x_4 - x_2 + x_3 - x_1)}. \tag{2}$$

This mixed strategy of non-monitoring P2 is non-degenerate. First, $q > 0$, given the payoff assumptions and $s < \frac{x_4-x_2}{x_4-x_1} < 1$. Next, $q < 1$ if and only if $s(x_3 - x_2) < x_3 - x_1$. If $x_3 - x_2 \leq 0$, this condition is trivially met. If $x_3 - x_2 > 0$, any $s < \frac{x_4-x_2}{x_4-x_1}$ meets the condition because $\frac{x_4-x_2}{x_4-x_1} < \frac{x_3-x_1}{x_3-x_2}$, or equivalently, $x_3 + x_4 > x_1 + x_2$. This condition is met due to the assumptions $x_3 > x_1$ and $x_4 > x_2$.

On the other hand, non-monitoring P2 uses a non-degenerate mixed strategy if and only if

$$py_4 + (1-p)y_1 = py_3 + (1-p)y_2, \tag{3}$$

which implies P1's probability of action A:

$$p = \frac{y_2 - y_1}{y_4 - y_3 + y_2 - y_1}. \tag{4}$$

Given the payoff assumptions, both $y_2 - y_1$ and $y_4 - y_3$ are strictly positive, ensuring that this mixed strategy of P1 is also non-degenerate.

The proof so far shows the existence of a unique equilibrium in each subgame following $s < \frac{x_4 - x_2}{x_4 - x_1}$. Note that q, the equilibrium strategy of non-monitoring P2 decreases in s while p does not. Intuitively, a high-probability of monitoring makes action A less appealing to P1. P1's indifference between the two actions require that P2 decreases her probability of action A when she does not monitor.

P2's ex-ante equilibrium payoff is

$$\begin{aligned}\Pi_2(s) &= s(py_4 + (1-p)y_2) + (1-s)(py_3 + (1-p)y_2) \\ &= y_2 + p(y_3 - y_2) + sp(y_4 - y_3) \\ &= \frac{y_4 y_2 - y_3 y_1}{y_4 - y_3 + y_2 - y_1} + s\frac{(y_4 - y_3)(y_2 - y_1)}{y_4 - y_3 + y_2 - y_1}.\end{aligned} \tag{5}$$

Given the assumptions on the payoffs, both $y_2 - y_1$ and $y_4 - y_3$ are strictly positive, ensuring that the coefficient for s is strictly positive. In other words, $\Pi_2(s)$ linearly increases in s.

(b) Note that q in part (a) is positive if and only if $s \leq \frac{x_4 - x_2}{x_4 - x_1}$. In particular, if $s > \frac{x_4 - x_2}{x_4 - x_1}$, then $sx_1 + (1-s)(qx_1 + (1-q)x_4) < sx_2 + (1-s)(qx_3 + (1-q)x_2)$ holds for all $q \in [0,1]$. Therefore, P1 playing action B with probability 1 and both monitoring and non-monitoring P2 doing the same constitutes the unique equilibrium for any $s > \frac{x_4 - x_2}{x_4 - x_1}$. In this case, $\Pi_1(s) = x_2$ and $\Pi_2(s) = y_2$.

(c) First, there is no equilibrium in which P1 takes action A with probability 1. If he did, and if P2 best responded by taking action A, P1's payoff would be x_1. Then P1 can strictly increase his payoff by deviating to action B, regardless of whether P2 observes his action of not.

Second, there exists an equilibrium in which P1 takes action B with probability 1. Given P2's best response, P1's payoff is x_2. A deviation to action A would yield $s^* x_1 + (1-s^*)x_4$, which is equal to x_2 given $s^* = \frac{x_4 - x_2}{x_4 - x_1}$. In this equilibrium, $\Pi_1(s) = x_2$ and $\Pi_2(s) = y_2$ as in (b).

Lastly, there are equlibria in which P1 plays a non-degenerate mixed strategy. Note that q in part (a) which induces P1's indifference between the two actions is zero, given that $s^* = \frac{x_4 - x_2}{x_4 - x_1}$. In other words, non-monitoring P2 takes action B with probability 1 in these equilibria. For action B to be incentive compatible for P2,

$$py_4 + (1-p)y_1 \leq py_3 + (1-p)y_2, \tag{6}$$

or equivalently,

$$p \leq \frac{y_2 - y_1}{y_4 - y_3 + y_2 - y_1}. \tag{7}$$

Therefore, in these equilibria, P1 takes action A with any probability $p \in [0, \frac{y_2-y_1}{y_4-y_3+y_2-y_1}]$, non-monitoring P2 takes action B with probability 1, and monitoring P2 matches the observed action by P1. Then P2's ex-ante equilibrium payoff is

$$\begin{aligned}
\Pi_2(s^*) &= s^*(py_4 + (1-p)y_2) + (1-s^*)(py_3 + (1-p)y_2) \\
&= y_2 + p(y_3 - y_2) + s^* p(y_4 - y_3) \\
&\in [y_2, y_2 + \frac{y_2 - y_1}{y_4 - y_3 + y_2 - y_1}(y_3 - y_2) + s^* \frac{y_2 - y_1}{y_4 - y_3 + y_2 - y_1}(y_4 - y_3)] \\
&\in [y_2, \frac{y_4 y_2 - y_3 y_1}{y_4 - y_3 + y_2 - y_1} + s^* \frac{(y_4 - y_3)(y_2 - y_1)}{y_4 - y_3 + y_2 - y_1}].
\end{aligned} \tag{8}$$

y_2 is P2's expected payoff in case (b), and $\frac{y_4 y_2 - y_3 y_1}{y_4-y_3+y_2-y_1} + s^* \frac{(y_4-y_3)(y_2-y_1)}{y_4-y_3+y_2-y_1}$ is the expected payoff she would get in case (a) if $s = s^*$.

A.2 Proof for Proposition

(a) First, the strategy profile described below is a perfect Bayesian equilibrium:

- P2 monitors with probability $s^* = \frac{x_4-x_2}{x_4-x_1} \in (0,1)$.
- If $s^* < \frac{x_4-x_2}{x_4-x_1}$, the players play according to the unique equilibrium prediction in Lemma (a).
- If $s^* > \frac{x_4-x_2}{x_4-x_1}$, the players play according to the unique equilibrium prediction in Lemma (b).
- If $s^* = \frac{x_4-x_2}{x_4-x_1}$, the players select the equilibrium described in the proof of Lemma (c) in which P1 takes action A with probability $p = \frac{y_2-y_1}{y_4-y_3+y_2-y_1}$.

Note that P2's ex-ante equilibrium payoff $\frac{y_4 y_2 - y_3 y_1}{y_4-y_3+y_2-y_1} + s^* \frac{(y_4-y_3)(y_2-y_1)}{y_4-y_3+y_2-y_1}$ and is strictly greater than y_2. To prove it,

$$
\begin{aligned}
s^* = \frac{x_4 - x_2}{x_4 - x_1} &> \frac{y_2 - y_3}{y_4 - y_3} \\
s^*(y_4 - y_3)(y_2 - y_1) &> (y_4 - y_3)(y_2 - y_1)\frac{y_2 - y_3}{y_4 - y_3} \\
s^*(y_4 - y_3)(y_2 - y_1) &> (y_4 - y_3 + y_2 - y_1)y_2 - (y_4y_2 - y_3y_1) \\
(y_4y_2 - y_3y_1) + s^*(y_4 - y_3)(y_2 - y_1) &> (y_4 - y_3 + y_2 - y_1)y_2 \\
\frac{y_4y_2 - y_3y_1}{y_4 - y_3 + y_2 - y_1} + s^*\frac{(y_4 - y_3)(y_2 - y_1)}{y_4 - y_3 + y_2 - y_1} &> y_2.
\end{aligned} \tag{9}
$$

Furthermore, P1's ex-ante equilibrium payoff is equal to x_2:

$$
\begin{aligned}
\Pi_1(s^*) &= s^*(x_1) + (1 - s^*)(qx_1 + (1 - q)x_4) \\
&= x_4 + q(x_1 - x_4) + s^*(1 - q)(x_1 - x_4) \\
&= x_4 + 0 \cdot (x_1 - x_4) + \frac{x_4 - x_2}{x_4 - x_1}(1 - 0)(x_1 - x_4) \\
&= x_2
\end{aligned} \tag{10}
$$

Therefore, Proposition (b) holds if (a) holds. Given Lemma (a) and (c), P2 also maximizes her ex-ante equilibrium payoff at $s^* = \frac{x_4 - x_2}{x_4 - x_1}$.

It remains to show the uniqueness of the equilibrium. First, there exists no equilibrium in which P2 chooses $s = s^*$ but the players selects a different equilibrium in the ensuing subgame. That would mean that P2 earns a strictly lower expected payoff than the one described above. Then, given Lemma (a), P2 can be strictly better off by deviating to a $s < s^*$ that is sufficiently close to s^*. For a similar reason, there exists no equilibrium in which P2 chooses $s < s^*$. In this case, P2 can be strictly better off by deviating to another $s < s^*$ that is closer to s^*. Lastly, there exists no equilibrium in which P2 chooses $s > s^*$. Given that $\frac{y_4y_2 - y_3y_1}{y_4 - y_3 + y_2 - y_1} + s^*\frac{(y_4 - y_3)(y_2 - y_1)}{y_4 - y_3 + y_2 - y_1}$ is strictly greater than y_2, there exists $s < s^*$ sufficiently close to s^* to which P2 can profitably deviate.

Appendix B Empirical optimality of over-monitoring: alternative definitions

Figure 7: P2's expected payoff conditional on probability of monitoring, assuming P1's empirical strategy and P2's best response

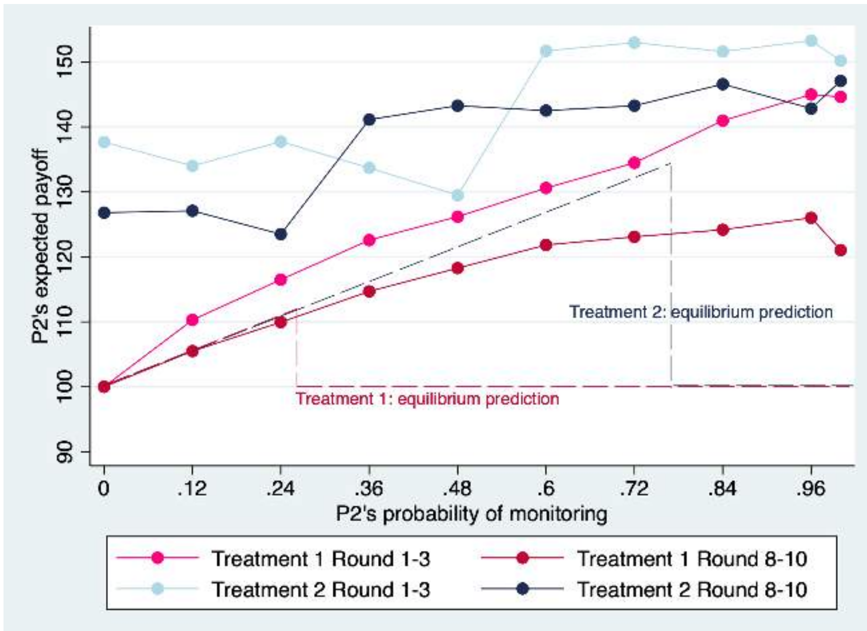

Figure 8: P2's mean empirical payoff conditional on probability of monitoring, given both players' empirical strategy.

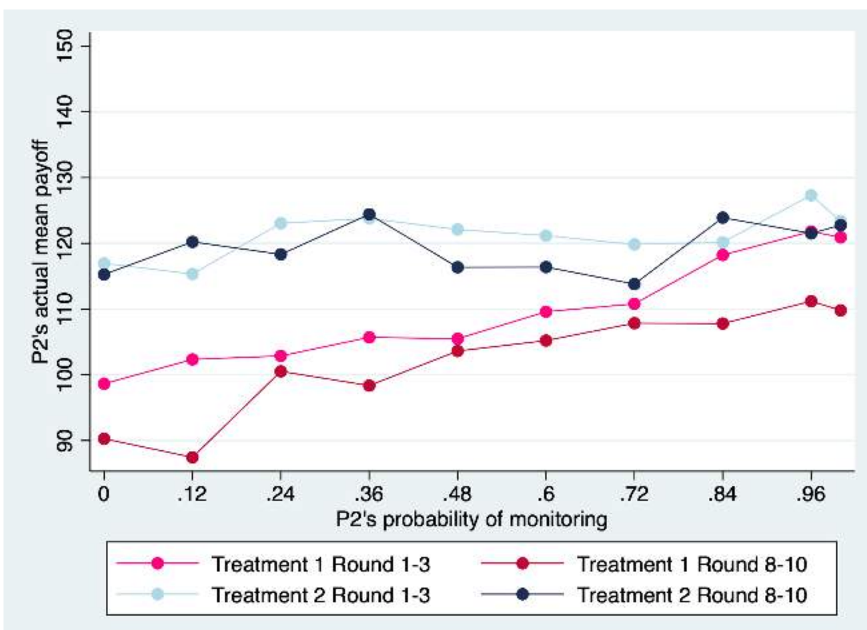

Appendix C P1's individual strategy

Figure 9: Individual P1's strategy, treatment 1

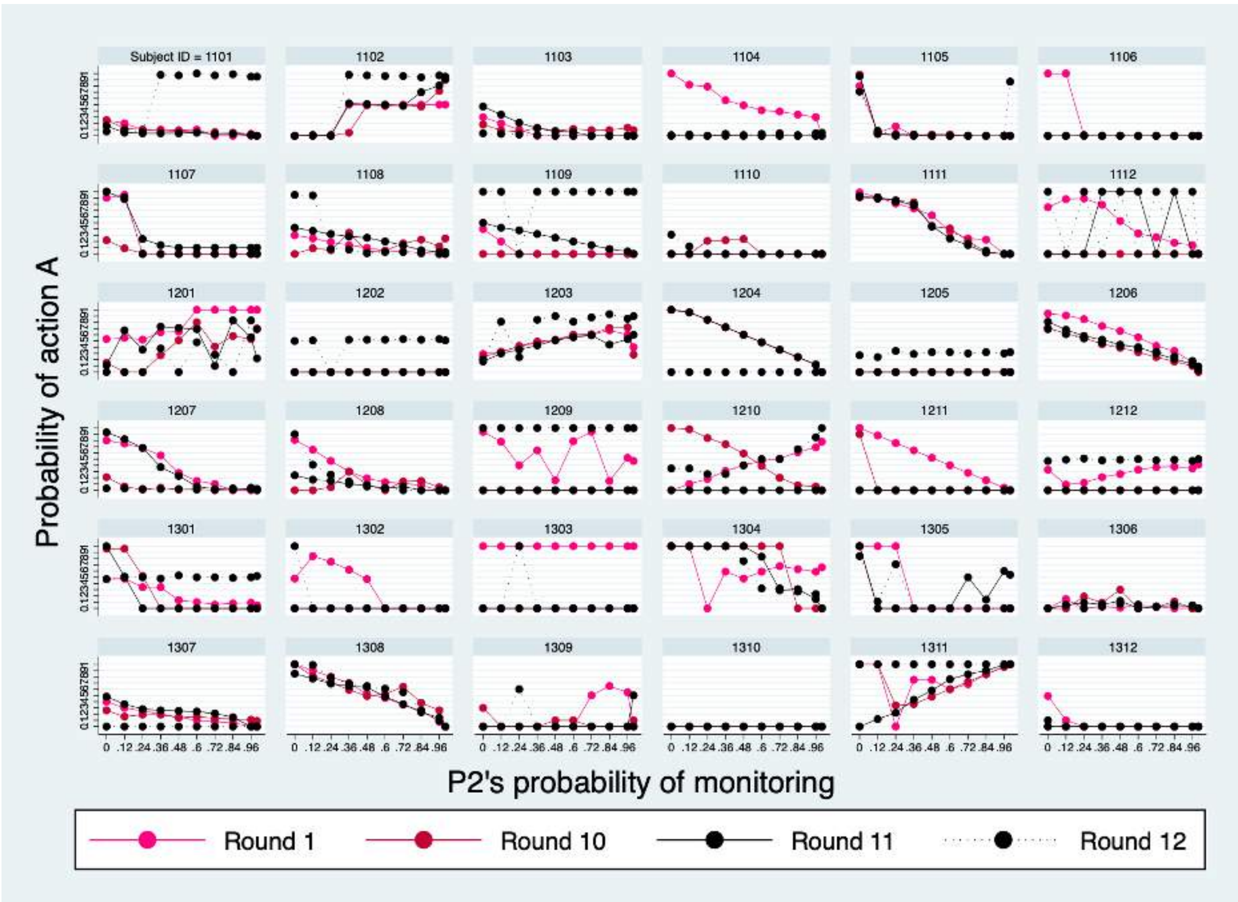

Figure 10: Individual P1's strategy, treatment 2

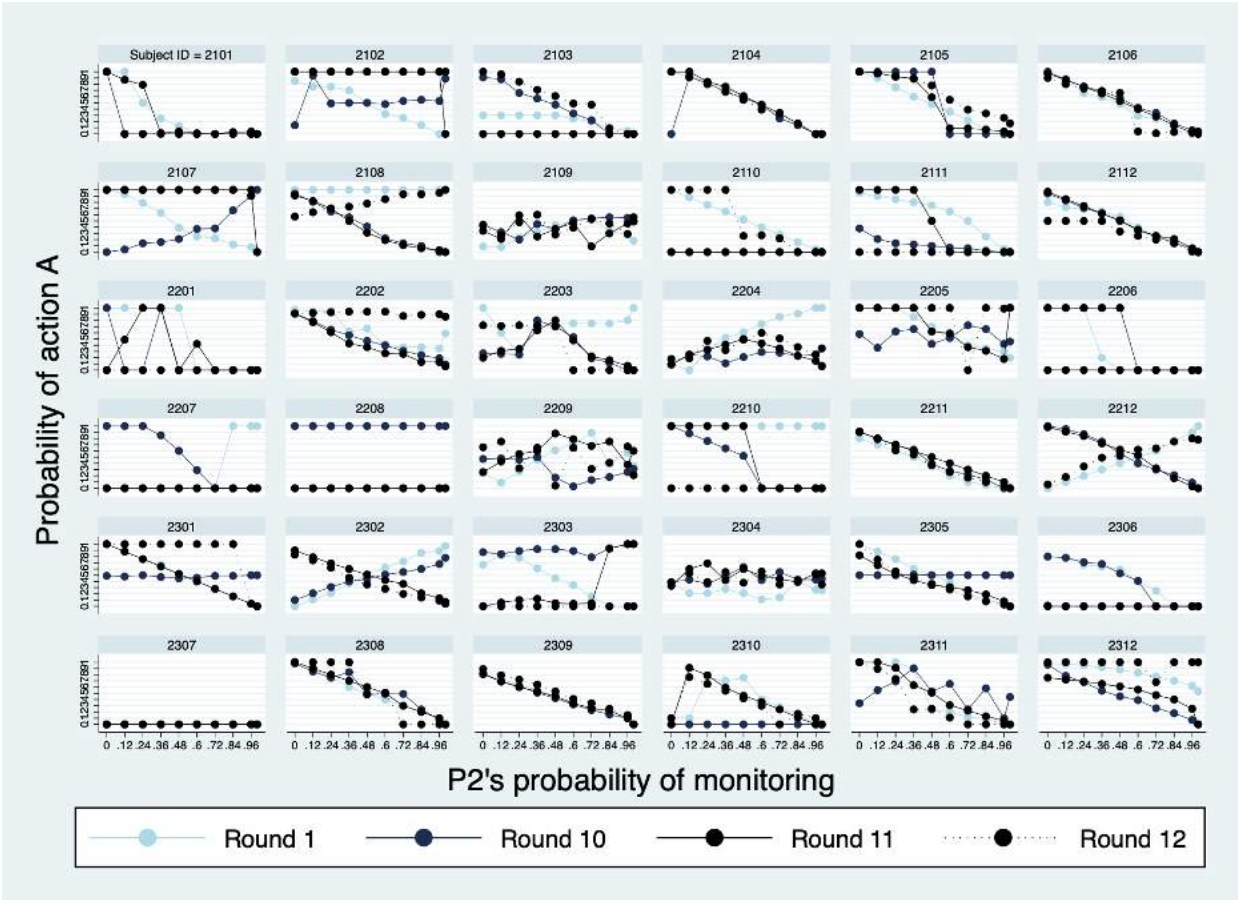

Appendix D Elicited beliefs: further analyses

D.1 Beliefs about P1's strategy

P2's first-order beliefs are largely consistent with P1's empirical strategy on average. Figure 11 shows P2's mean first-order belief and P1's mean second-order belief about P1's probability of action A in treatment 1. Figure 12 shows the corresponding beliefs in treatment 2. Specifically, in both treatments, P2's believe that P1's probability of action A decreases in P2's probability of monitoring (OLS regression of belief in round 10 on probability of monitoring, clustered error at the subject level, $\beta = -.1860, p = 0.005$ in treatment 1; $\beta = -.2931, p = 0.000$ in treatment 2). P1's second-order beliefs show a similar pattern, although it is not statistically significant only in treatment 1 ($\beta = -.1102, p = 0.135$ in treatment 1; $\beta = -.2472, p = 0.021$ in treatment 2).

Comparing the beliefs between the two treatments help us understand the over-monitoring observed in treatment 1. For example, compared to treatment 2, P2's in treatment 1 are less likely to hold first-order beliefs that rationalizes low-probability monitoring. Figures 13 and 14 in the appendix show individual P2's first-order beliefs about P1's strategy. The figures show that P2's in treatment 1 are less likely to believe that P1's use strategies similar to those predicted by the equilibrium (DDS as described in section 5.2.3). 10 out of 52 P2's in treatment 1, and 17 out of 53 P2's in treatment 2 hold such beliefs in round 10. Furthermore, such beliefs correlate to low-probability monitoring in treatment 2 (two-sided ranksum test, $p \leq 0.0368$ in round 10 and 11), but not in treatment 1 ($p \geq 0.6437$ in round 1, 10, and 11).

P1's second-order beliefs also fit into a coherent picture. Figures 15 and16 show individual P1's second-order beliefs about P2's first-order beliefs. In treatment 1, P1's often report belief that is constant across the probability of monitoring. Few P1's in treatment 1 (2 out of 36 in round 10) believe that P2's first-order beliefs are consistent with DDS, whereas a substantial number of P1's in treatment 2 (12 out of 36 in round 10) hold such beliefs. Furthermore, P1'second-order beliefs in DDS correlate to actual use of DDS in treatment 2 (two-sided ranksum test, $p \leq 0.0288$ in round 1 and 11; $p = 0.3404$ in round 10).

Figure 11: P2's first-order belief & P1's second-order belief about P1's strategy, treatment 1

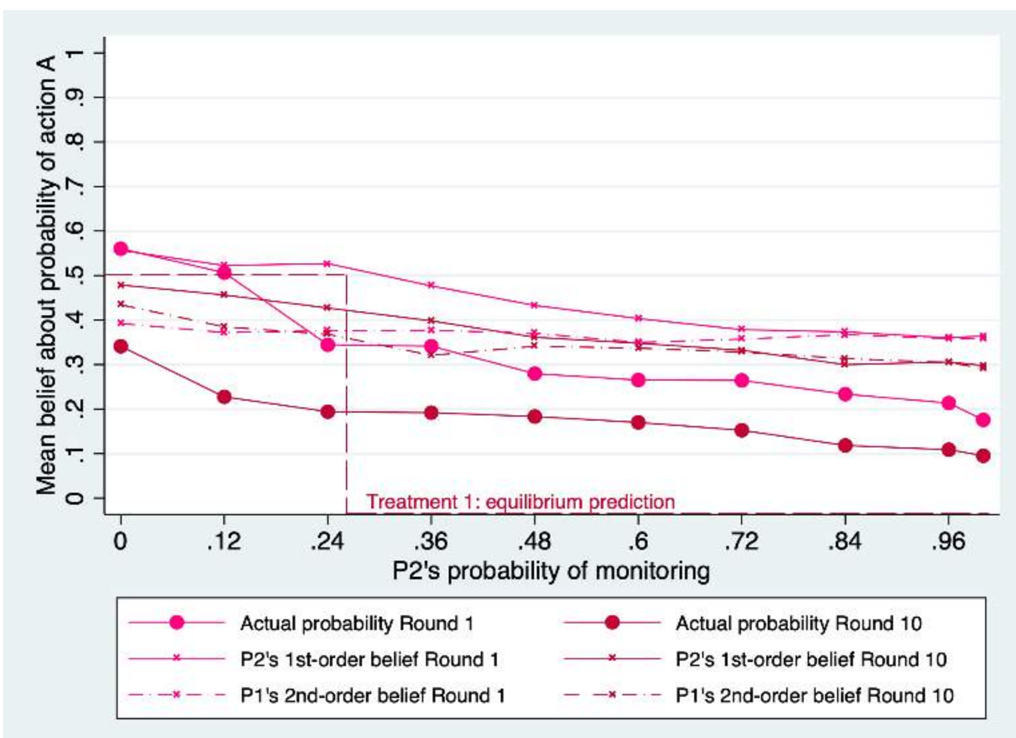

Figure 12: P2′s first-order belief & P1′s second-order belief about P1′s strategy, treatment 2

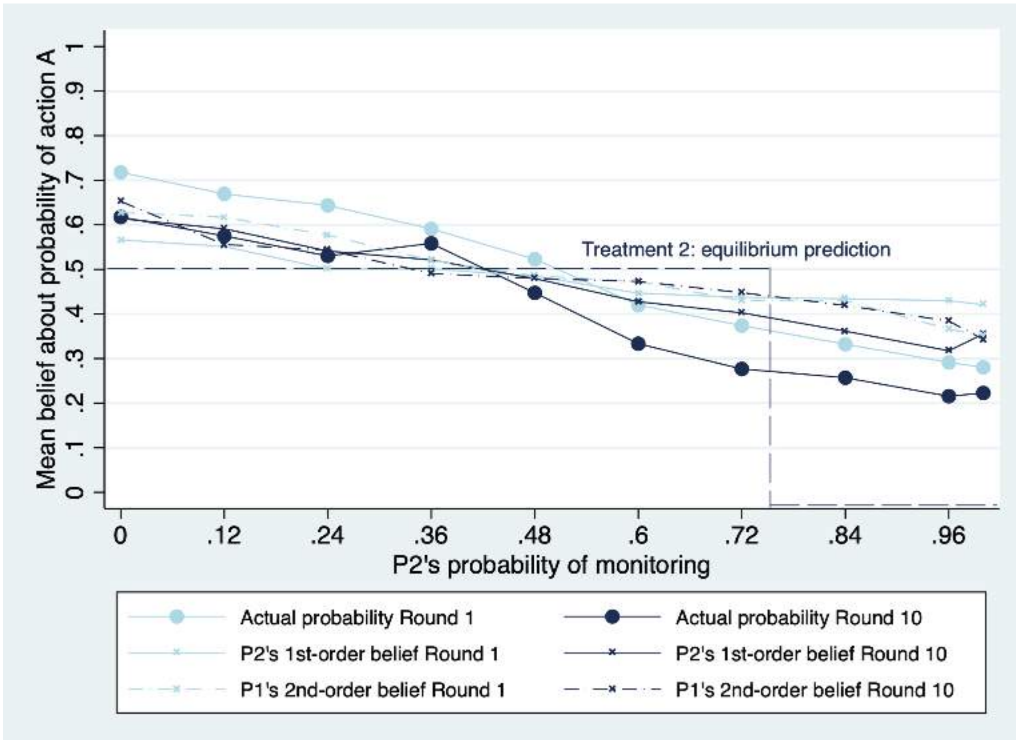

Figure 13: Individual belief about P1's strategy: P2's first-order belief, treatment 1

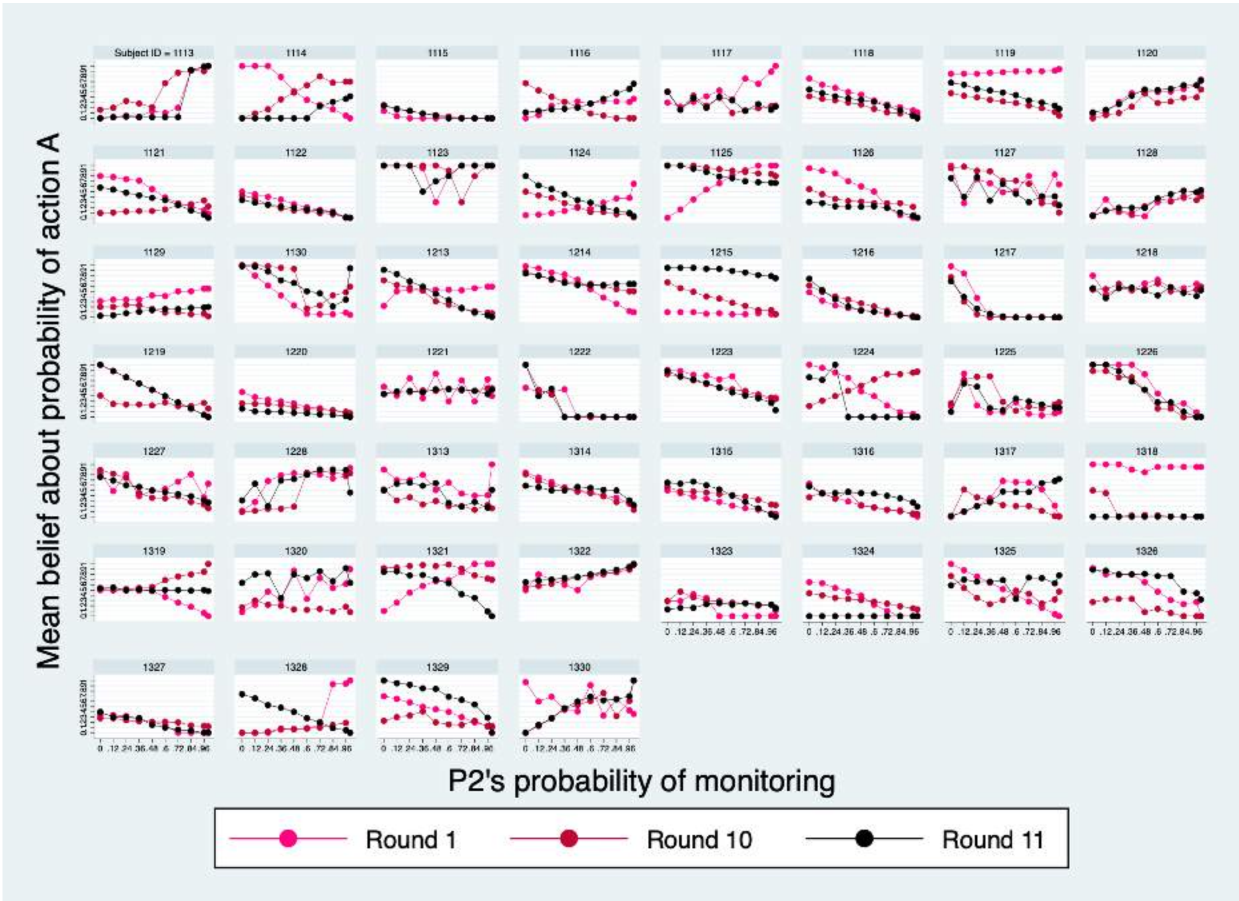

Figure 14: Individual belief about P1′s strategy: P2′s first-order belief, treatment 2

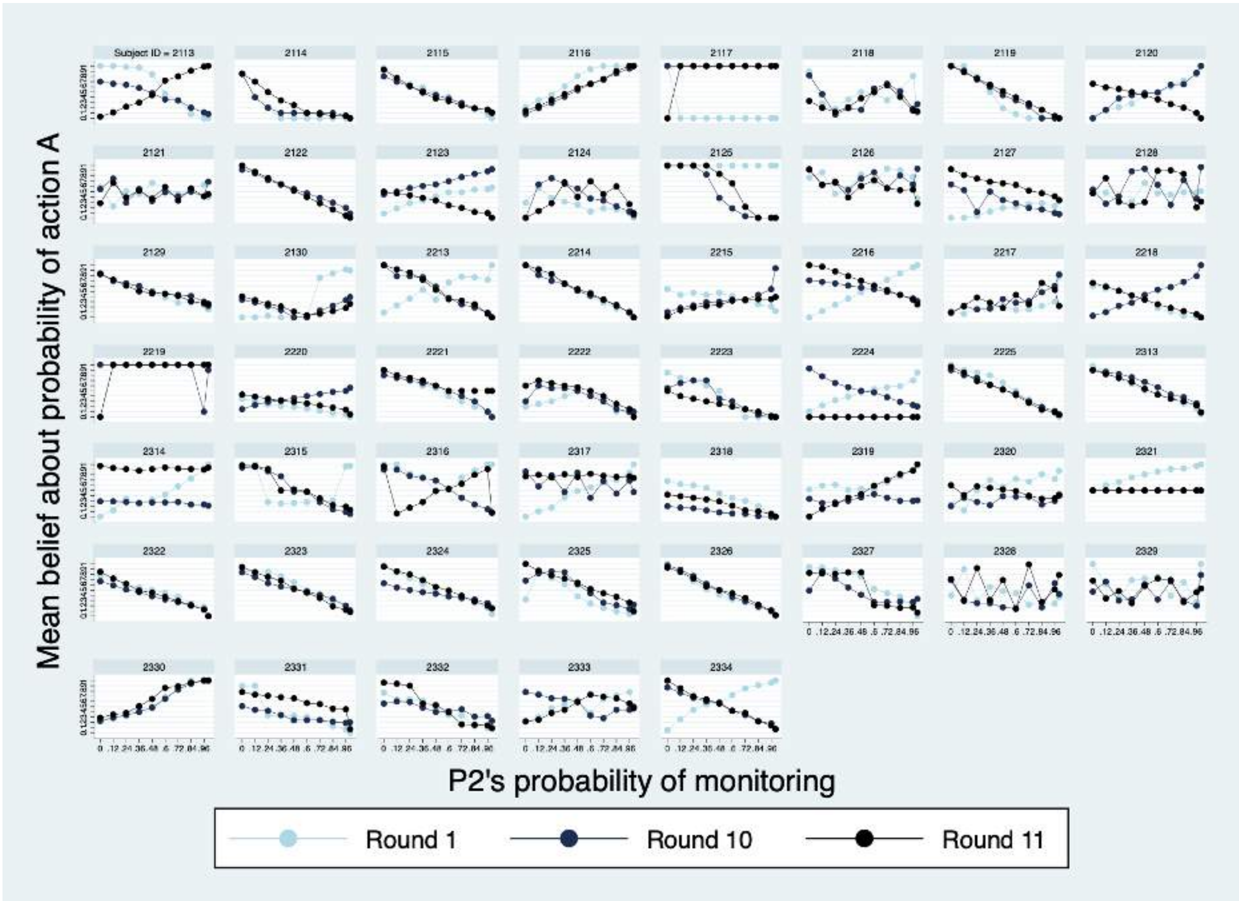

Figure 15: Individual belief about P1's strategy: P1's second-order belief, treatment 1

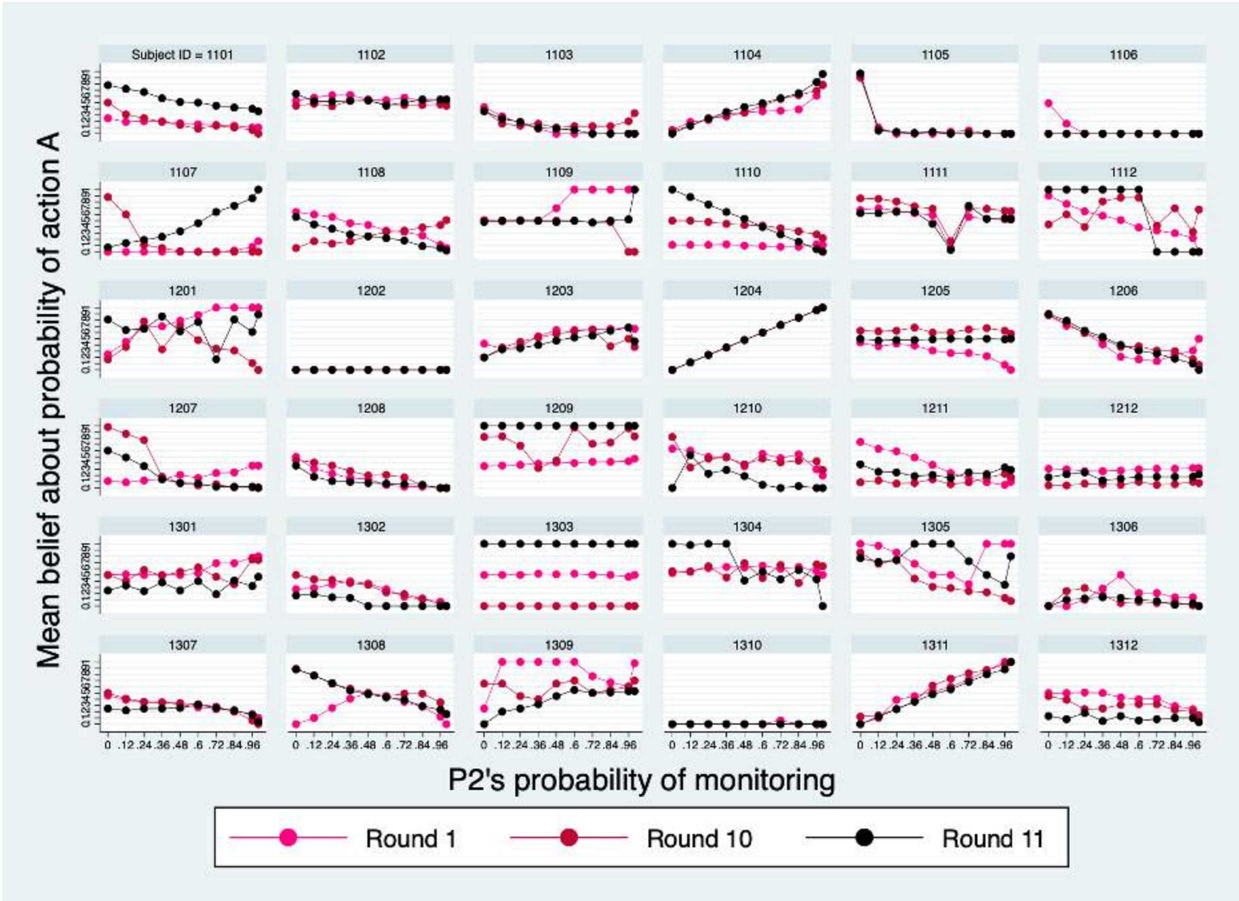

Figure 16: Individual belief about P1's strategy: P1's second-order belief, treatment 2

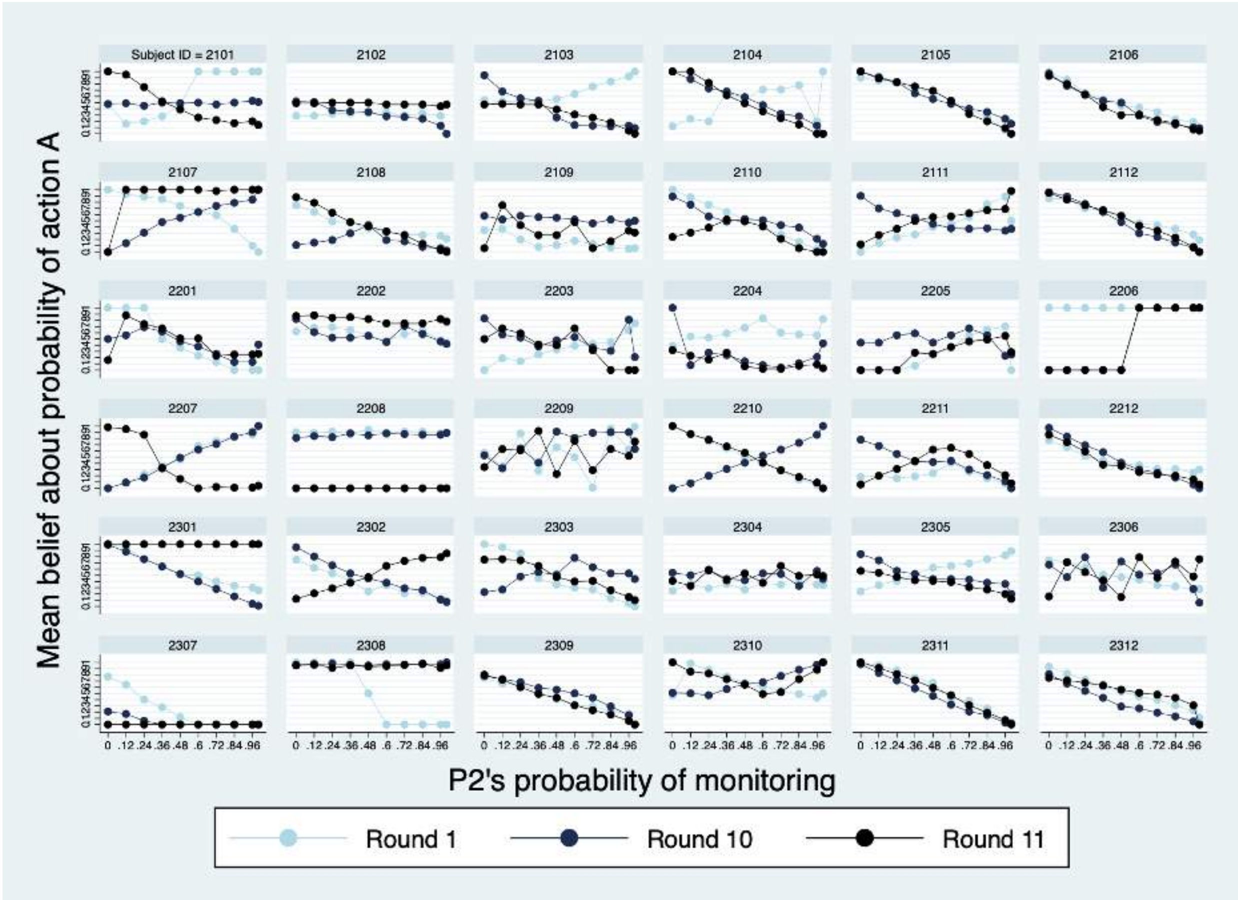

D.2 Beliefs about P2's strategy

Figures 17 and 18 show the mean beliefs about non-monitoring P2's strategy in each treatment. Unlike the beliefs about P1's strategy that assume a negative relationship between the monitoring probability and action A ("innovate"), the beliefs about P2's strategy are constant across the monitoring probabilities in both treatments (OLS regression of belief in round 10 on probability of monitoring, clustered error at the subject level; P1's mean first-order beliefs $\beta = .0591527, p = 0.417$ in treatment 1; $\beta = -.0054354, p = 0.962$ in treatment 2; P2's mean second-order beliefs $\beta = -.0714784, p = 0.255$ in treatment 1; $\beta = .007024, p = 0.924$ in treatment 2). Furthermore, the mean beliefs fail to reflect the fact that P2's strategy converges towards its respective equilibrium prediction in each treatment.

The mean beliefs are similar between the two treatments yet they have opposite implications for subjects' behavior. The equilibrium predictions, as shown in dashed lines in figures 17 and 18, requires P2 to play action A with much lower probability in treatment 1

Figure 17: Mean belief about non-monitoring P2′s strategy: P1′s first-order belief and P2′s second-order belief, treatment 1

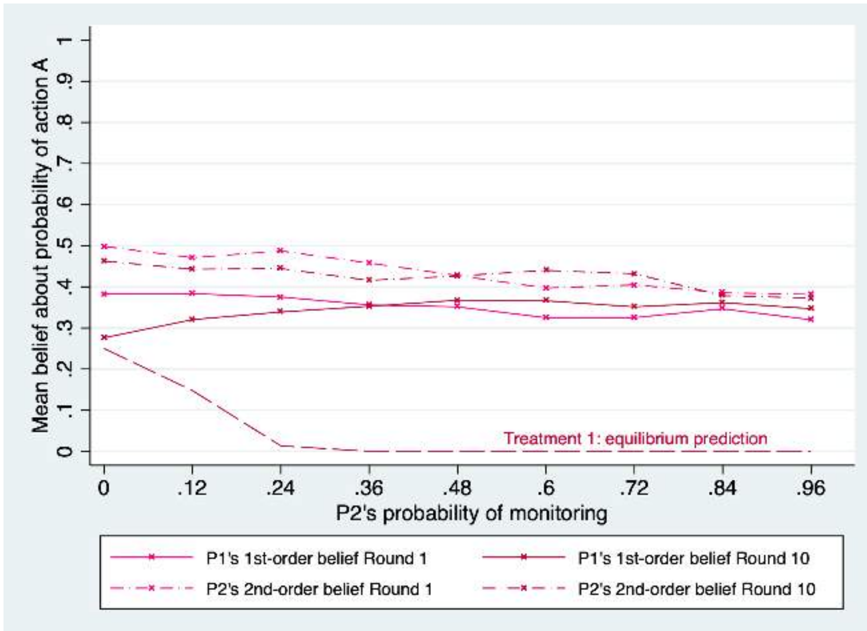

Figure 18: Mean belief about non-monitoring P2′s strategy: P1′s first-order belief and P2′s second-order belief, treatment 2

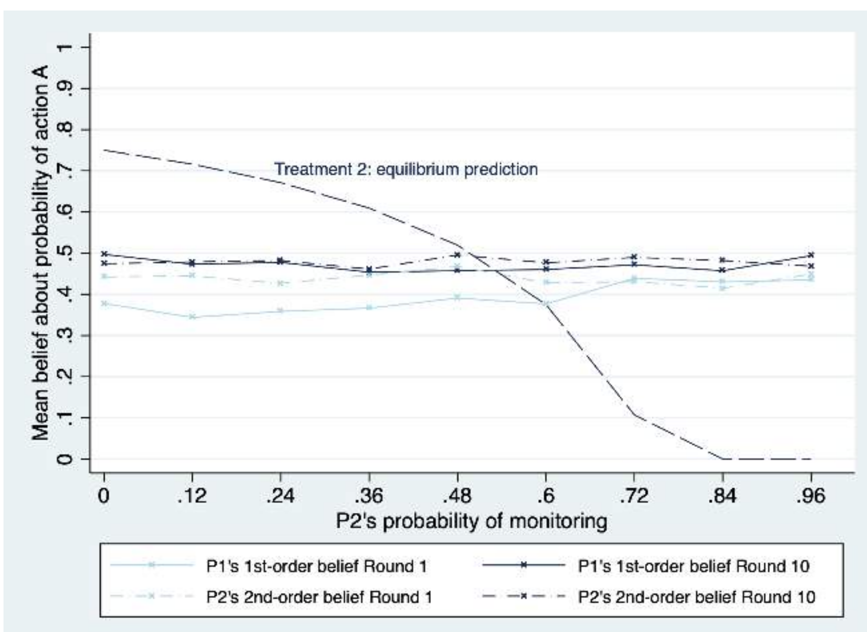

compared to treatment 2, to preserve P1's incentive to play action A ("innovate"). Therefore, a sequentially rational P1 who holds the mean beliefs would play action A under the low probability of monitoring in treatment 2, whereas he would not in treatment 1 even under a low probability of monitoring.

Similar observation can be made from the individual beliefs shown in figures 19 and 20 in the appendix. In treatment 1, a majority of P1's believe that the non-monitoring P2 overplay action A conditional on $s < s^* = 0.25$ (64.81%, 70 out of 108 observations in round 10). In treatment 2, less than half of P1's hold such beliefs conditional on $s < s^* = 0.75$ (43.65 %, 110 out of 252 observations in round 10). The data also suggest that P1's who hold such beliefs are less likely to take action A under low probability of monitoring ($s < s^*$) in both treatments. (OLS regression of probability of action A on dummy variable for the belief that non-monitoring P2 overplays A, round 10; clustered error at the subject level; $\beta = -.1911579, p = 0.088$ in treatment 1; $\beta = -.1338067, p = 0.077$ in treatment 2)

P2's mean second-order beliefs closely approximates P1's mean first-order beliefs, except under the low probabilities of monitoring in treatment 1. In those exceptional cases, P2's slightly overestimate P1's first-order beliefs (two-sided ranksum test, $p \leq 0.0416$ for each $s < s^* = 0.25$ in treatment 1, round 10: $p \geq 0.1575$ for each $s > s^* = 0.25$ in treatment 1, round 10). In those cases, most P2's also believe that P1's anticipate the non-monitoring P2's to overplay A (82.69%, 129 out of 156 observations, $s < s^* = 0.25$ in treatment 1, round 10).

The second-order beliefs of P2's imply that they would hold relatively pessimistic first-order beliefs about P1's chance of playing action A in treatment 1. Such pessimistic beliefs, which were confirmed in section D.1, would imply action B under sequential rationality. However, Figure 6 in Section 5.3 showed that P2's tend to overplay action A in treatment 1. Further analyses show that P2's second-order belief is not a statistically significant predictor of her monitoring behavior or the probability of action A:

- OLS regression of monitoring probability on second-order beliefs conditional on $s < s^*$; clustered error at the subject level; $\beta = -.0676, p = 0.686$ in treatment 1, round 10; $\beta = -.0559, p = 0.766$ in treatment 2, round 10

- OLS regression of monitoring probability on dummy variable for second-order belief of overplayed action A conditional on $s > s^*$; clustered error at the subject level; $\beta = .0007, p = 0.992$ in treatment 1, round 10; $\beta = -.0099, p = 0.888$ in treatment 2, round 10

- OLS regression of probability or action A on second-order beliefs conditional on $s > s^*$; clustered error at the subject level; $\beta = -.0676, p = 0.686$ in treatment 1, round 10; $\beta = -.0559$, $p = 0.766$ in treatment 2, round 10

- OLS regression of monitoring probability on dummy variable for second-order belief of overplayed action A conditional on $s > s^*$; clustered error at the subject level; $\beta = .0007$, $p = 0.992$ in treatment 1, round 10; $\beta = -.0099$, $p = 0.888$ in treatment 2, round 10

Figure 19: Individual belief about non-monitoring P2's strategy: P1's first-order belief, treatment 1

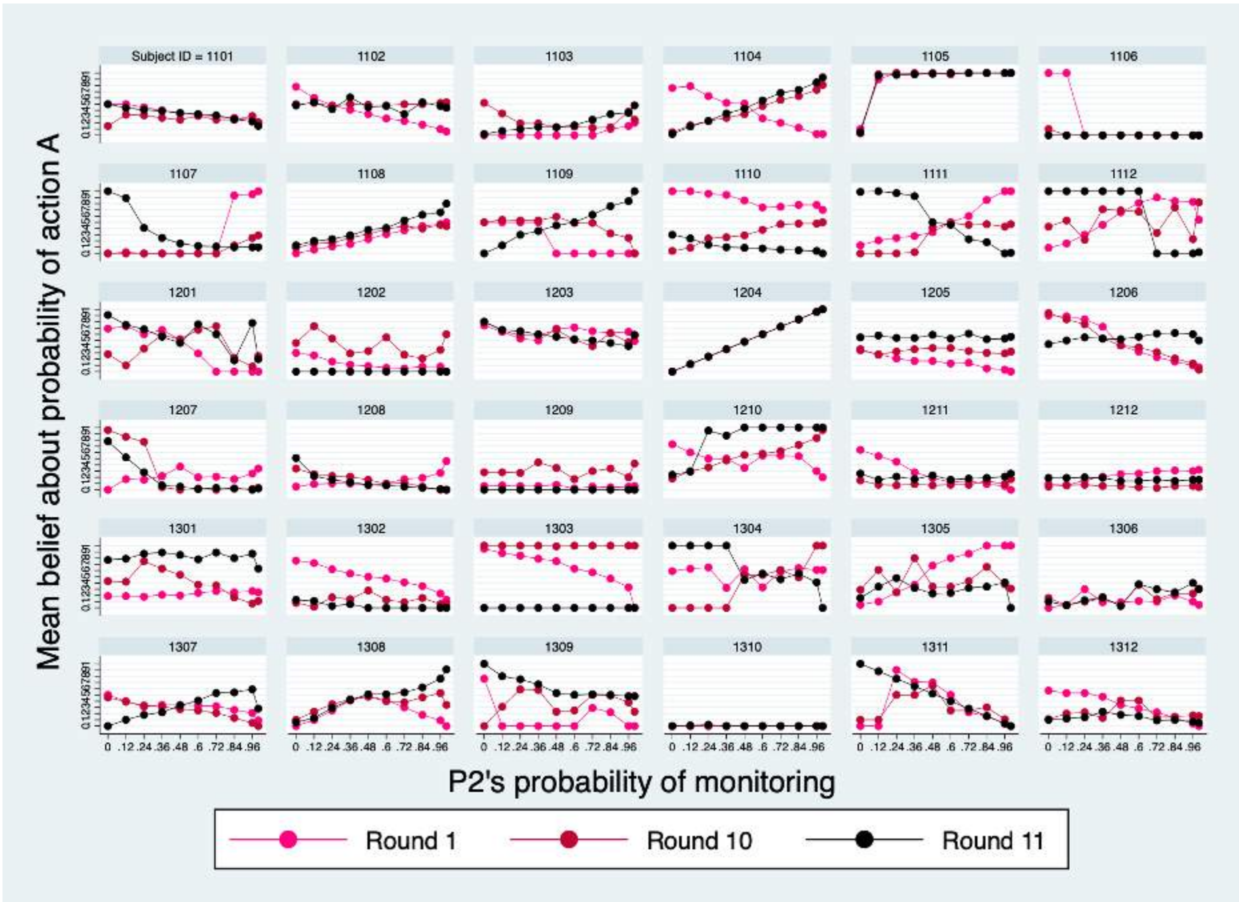

Figure 20: Individual belief about non-monitoring P2's strategy: P1's first-order belief, treatment 2

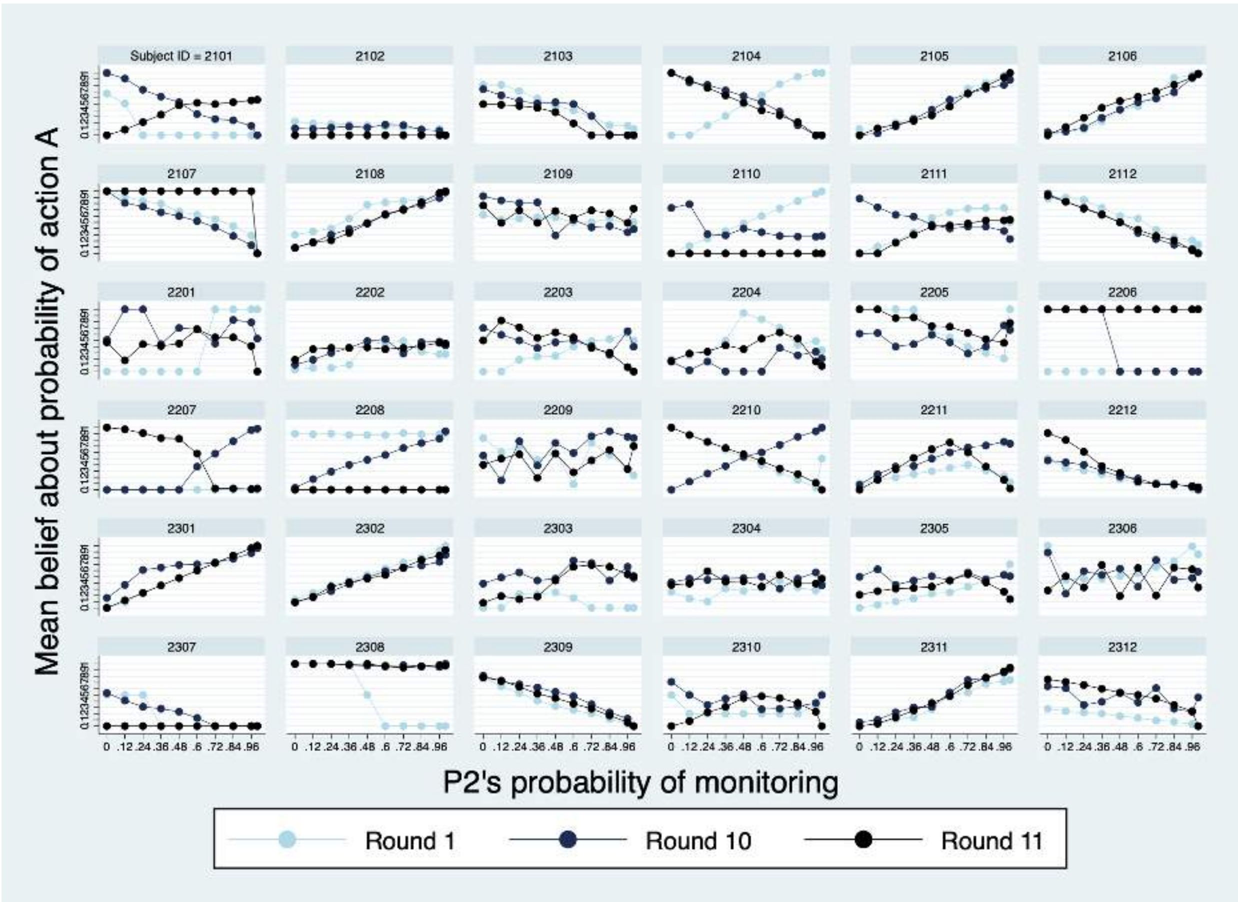

Figure 21: Individual belief about non-monitoring P2's strategy: P2's second-order belief, treatment 1

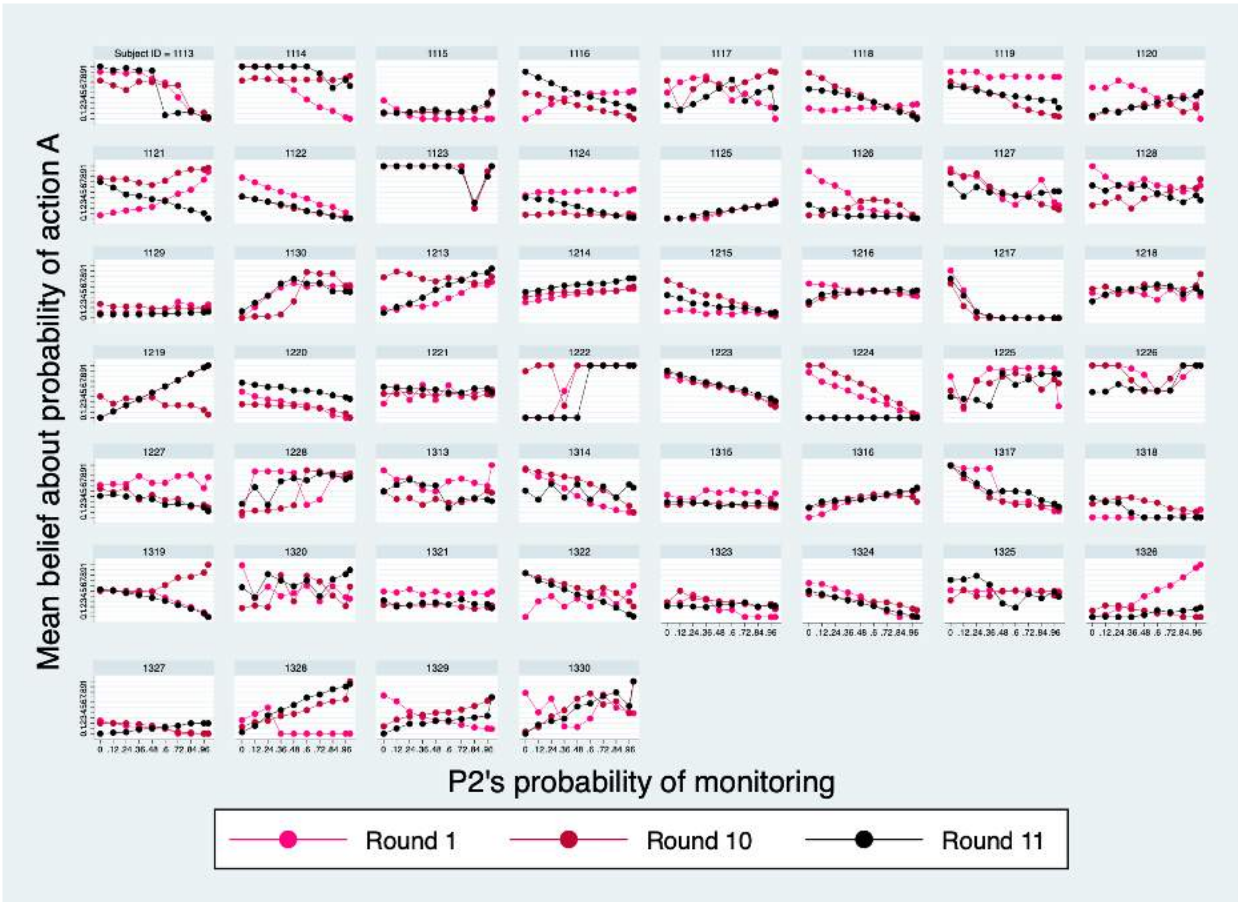

Figure 22: Individual belief about non-monitoring P2's strategy: P2's second-order belief, treatment 2