

Sentiments, strategic uncertainty, and information structures in coordination games*

Michal Szkup[†] and Isabel Trevino[‡]

Abstract

We study experimentally how changes in the information structure affect behavior in coordination games with incomplete information (global games). We find two systematic departures from the theory: (1) the comparative statics of equilibrium thresholds and signal precision are reversed, and (2) as information becomes very precise subjects' behavior approximates the efficient equilibrium of the game, not the risk dominant one. To organize our findings we extend the standard global game model to allow for sentiments in the perception of strategic uncertainty and study how they relate to fundamental uncertainty. We test the extended model by eliciting first-order and second-order beliefs and find support for the sentiments mechanism: subjects are over-optimistic about the actions of others when the signal precision is high and over-pessimistic when it is low. Thus, we show how changes in the information structure can give rise to sentiments that drastically affect outcomes in coordination games. This novel mechanism can help explain stylized facts and offer policy guidance for environments characterized by strategic complementarities and incomplete information.

Keywords: global games, coordination, information structures, strategic uncertainty, sentiments, biased beliefs.

JEL codes: C72, C9, D82, D9

1 Introduction

Many economic phenomena can be analyzed as coordination problems under uncertainty. Investment decisions, currency attacks, bank runs, or political revolts illustrate situations where decision makers would like to coordinate with others to attain certain outcomes but may fail to do so. In

*Financial support from the National Science Foundation via grant SES-105962 and the Center for Experimental Social Science at New York University is greatly appreciated. This paper incorporates some results from our working paper entitled “Costly information acquisition in a speculative attack: Theory and experiments” and supersedes it. All remaining errors are ours.

[†]Vancouver School of Economics, University of British Columbia, 6000 Iona Drive, Vancouver, BC, V6T 1L4, Canada, michal.szkup@ubc.ca.

[‡]Department of Economics, University of California San Diego, 9500 Gilman Drive #0508 La Jolla, CA 92093, USA. itrevino@ucsd.edu.

addition to the strategic uncertainty that arises from not knowing the actions of others, in these environments decision makers also face uncertainty about the fundamentals that determine the state of the economy (e.g., the profitability of the investment, the strength of the currency peg, or the strength of the political regime). In these environments the information structure characterizes both the degree of fundamental and strategic uncertainty faced by the players, and hence determines the coordination outcome. Thus, a key question is whether better information leads to less coordination failure (see e.g., Angeletos and Lian (2016), Veldkamp (2011), or Vives (2010)).

In this paper we investigate theoretically and experimentally how changes in the information structure affect behavior in coordination games with incomplete information. We use global games as the setup to perform our analysis because they offer a unique framework to study explicitly the effects of fundamental and strategic uncertainty on behavior (see Carlsson and Van Damme (1993), Morris and Shin (1998, 2003)).¹ Global games are coordination games with incomplete information where players observe noisy private signals about payoffs. This perturbation in the information structure leads to a unique equilibrium under mild conditions on parameters and is characterized by coordination failures.² In these games, the precision of the signals determines the degree of fundamental and strategic uncertainty. This leads to sharp theoretical predictions: in two-player settings, as the signal precision increases (fundamental uncertainty decreases), equilibrium play is driven more by strategic uncertainty and less by fundamental uncertainty.

Our experimental results show two departures from the theoretical predictions of global games: (1) the comparative statics of thresholds with respect to signal precisions are reversed and (2) as the signal noise decreases, subjects' behavior tends towards the efficient threshold, not the risk-dominant one. These results are robust to a number of variations in the experimental setting. We argue and provide evidence that these departures are driven by sentiments (i.e., biased beliefs) in the perception of strategic uncertainty and that these sentiments are linked in a systematic way to the degree of fundamental uncertainty. Thus, we identify behavioral biases that may crucially affect outcomes in strategic environments with coordination motives.

We begin the paper by setting up a standard global game model that we use as a theoretical benchmark for our experiment. We characterize its unique equilibrium and comparative statics and explain how the equilibrium is affected by the degree of fundamental and strategic uncertainty. In the experiment we vary fundamental uncertainty exogenously (exogenous variations of signal precision) and endogenously (costly information acquisition).³ We find that in both settings the

¹Global games are coordination games with incomplete information where players observe noisy private signals about payoffs.

²Global games have been used extensively in applications, such as Angeletos et al. (2006), Edmond (2013), Goldstein and Pauzner (2005), Hellwig et al. (2006), or Szkup (2016). See Angeletos and Lian (2016) for an excellent summary of this literature.

³We endogenize the information structure to study an environment where subjects can control the degree of fundamental uncertainty in an effort to better understand the relationship between fundamental and strategic uncertainty.

vast majority of subjects use threshold strategies, as suggested by the theory and consistent with existing experimental evidence (see Heinemann, Nagel, and Ockenfels (2004, 2009)). However, as mentioned above, we find two systematic departures in the way that subjects respond to changes in the information structure, which have significant welfare effects and are more pronounced when the information structure is endogenously determined by subjects. These departures also suggest that, contrary to the theoretical predictions, the perception of strategic uncertainty of might be directly aligned to the degree of fundamental uncertainty in the environment. This is because subjects behave as if they were more certain about the action of their opponent when signals become more precise.

To reconcile our findings with the theory, we extend the standard model of global games to allow for situations where the perception of strategic uncertainty can be influenced by sentiments. This extension is based on the observation that the expected payoffs in our model can be approximated by the product of the expected value of the state (determined by the degree of fundamental uncertainty) and the probability that the other player takes the risky action (which captures strategic uncertainty). This simple observation suggests that sentiments can be related to either fundamental uncertainty (biased beliefs about the state of fundamentals) or to strategic uncertainty (biased beliefs about the likelihood of the opponent taking a specific action).⁴ However, when information is very precise, signals convey very accurate information about the state, suggesting that, at least in the case of high precision, departures from the theory are likely to be driven by a biased perception of strategic uncertainty. Therefore, we hypothesize that our experimental findings are a result of sentiments that affect the perception of strategic uncertainty. Moreover, we hypothesize that the nature of these sentiments (i.e., their sign and magnitude) is directly related to the degree of fundamental uncertainty in the environment.

We test the sentiment-based mechanism of the extended model by eliciting subjects' first- and second-order beliefs and we find support for our hypotheses. On average, subjects form accurate first-order beliefs, especially with high and medium precisions, supporting the idea that sentiments related to fundamentals are an unlikely driver of our results. On the other hand, elicited second-order beliefs indicate that subjects are overly optimistic about the desire of their opponent to coordinate when information is very precise, and pessimistic when the signal noise increases. This suggests that subjects anchor their perception of strategic uncertainty to the degree of fundamental uncertainty, that is, fundamental uncertainty determines the sign and magnitude of the sentiments that affect the subjects' perception of strategic uncertainty.

The results of our experiments show how a bias in belief formation that has been extensively studied in individual decision making can crucially affect outcomes in strategic environments by

⁴The biased beliefs related to strategic uncertainty can be due, for example, to a player believing that her opponent has a biased perception of fundamental uncertainty.

altering the way in which players respond to changes in the information structure. The sentiments that we identify, however, are different to the biases that are typically studied in the behavioral literature (which focuses mainly on individual decision making) because they affect the perception of strategic uncertainty, which is unique to strategic environments.⁵ We find that the sign and magnitude of the bias we identify depend on the informativeness of the environment, which has important welfare implications for coordination problems. For example, more successful coordination can be attained when information is very precise because it leads players to become overoptimistic about the likelihood of a successful coordination. On the other hand, coordination failures are likely to occur under very noisy information because players tend to become overly pessimistic about the likelihood of a successful coordination.⁶

Our results not only provide novel insights about behavior in coordination problems with incomplete information, but they also shed light on some of the recent findings in macroeconomics and finance. For example, in the context of business cycles, Bloom (2009) suggests that recessions are accompanied by an increase in uncertainty, while Angeletos and La'O (2013) and Benhabib et al (2015) argue that recessions can be driven by sentiments. Our framework provides a natural connection between these two seemingly unrelated ideas. As our results show, an increase in uncertainty leads to negative sentiments about the likelihood of profitable risky investment (via pessimism about others investing). This leads to lower levels of aggregate investment, which leads to, or amplifies, a recession. Our results can also inform policy making. For example, the mechanism we identify can help to make the case for greater transparency in financial regulation. If we interpret rolling over loans as a risky choice in environments with strategic complementarities (as in Diamond and Dybvig (1983) or Goldstein and Pauzner (2005)), an increase in transparency (better information) can lead to positive sentiments about the likelihood of others rolling over, which decreases coordination failure by having fewer early withdrawals and can result in greater financial stability.

Our setup with exogenous information structures is a discrete version of Morris and Shin (1998) and the corresponding experimental treatments are related to Heinemann, Nagel, and Ockenfels (2004) who test the predictions of Morris and Shin (1998). The setup with endogenous information structures is related to Szkup and Trevino (2015) and Yang (2015) who endogenize the information structure in global games by allowing players to choose the quality of their information, at a cost.

⁵If the sentiments were related to the perception of fundamental uncertainty (beliefs about the state), they would be closer to the biases studied in the individual decision making literature (see Weinstein (1980) or Camerer and Lovo (1999)).

⁶Even though the bias we identify is related to strategic and not fundamental uncertainty, our characterization is qualitatively consistent with the characterization of the related biases in the psychology literature. For example, Moore and Cain (2007) show that subjects tend to be more pessimistic/underconfident when tasks are more difficult, and that they become optimistic/overconfident for easier tasks. These results are consistent with our characterization if we interpret the level of uncertainty as determining the level of difficulty to coordinate in a game.

Our model with sentiments is related to Izmalkov and Yildiz (2010) who study theoretically how sentiments affect outcomes in global games of regime change. However, we use a different notion of sentiments which is driven by our experimental findings.⁷

Our paper contributes to the literature that studies coordination games both experimentally and theoretically. Harsanyi and Selten (1988) define risk dominance and payoff dominance as two contrasting equilibrium refinements for coordination games with multiple equilibria. They suggest that in the presence of Pareto ranked equilibria risk dominance is irrelevant since “collective rationality” should select the payoff dominant equilibrium. However, experimental evidence highlights how strategic uncertainty can lead to coordination failure in games with complete information and Pareto ranked equilibria (see Van Huyck, Battalio, and Beil (1990, 1991), Cooper, DeJong, Forsythe, and Ross (1990, 1992), or Straub (1995)). More recent literature studies how the information structure affects equilibrium in coordination games. Theoretical contributions include Angeletos and Pavan (2007), Bannier and Heinemann (2005), Colombo, Femminis and Pavan (2014), Hellwig and Veldkamp (2009), Iachan and Nenov (2015), Pavan (2016), Szkup and Trevino (2015), and Yang (2015), among others. Darai, Kogan, Kwasnica, and Weber (2017) and Avoyan (2017) use a global game setting to study experimentally how different types of public signals and cheap talk affect coordination outcomes, respectively. Cornand and Heinemann (2014) and Baeriswyl and Cornand (2016) propose alternative ways of thinking about coordination in experiments about the closely related family of beauty contest games. This paper also contributes to the literature that incorporates aspects of bounded rationality to propose alternative equilibrium notions such as Nagel (1995), Costa-Gomes and Crawford (2006), McKelvey and Palfrey (1995), or Eyster and Rabin (2005).

The paper is structured as follows. Section 2 presents the theoretical benchmarks for the treatments in the experiment. Section 3 presents the experimental design and the theoretical predictions for the parameters used in the experiment. In Section 4 we present our experimental findings and characterize the main departures from the theory. In Section 5 we propose an extension to the model of Section 2 that allows for sentiments in the perception of strategic uncertainty in an effort to reconcile our findings with the theory and we provide evidence to support it. Section 5.4 discusses alternative explanations for the observed departures from the theory and concludes.

2 The model

In this section, we describe the theoretical models that serve as benchmarks for our experiment. We first describe the general version of our model where signals have heterogenous precisions across

⁷In Izmalkov and Yildiz (2010) sentiments arise because of a non-common prior assumption. In our setup all subjects share the same prior belief and sentiments arise because of the uncertainty that a player faces regarding the interpretation and/or use of information by his opponent.

players. A special case of this model, where precision is homogenous across players, corresponds to the standard global game with exogenous information structures, as in Carlsson and van Damme (1993) and Morris and Shin (2003). We then consider the case with endogenous information structures that is composed of a first stage of costly information acquisition and a second stage where players play the game with heterogeneous precisions.⁸

2.1 The setup

There are two identical players in the economy, $i \in \{1, 2\}$, who simultaneously choose whether to take action A or action B . Action B is safe and always delivers a payoff of 0. Action A is risky and has a cost T associated to it. Action A delivers a payoff of $\theta - T$ if it is successful and $-T$ if it fails, where $\theta \in \mathbb{R}$ is a random variable that reflects the state of the economy. Action A succeeds if both players choose action A and $\theta > \underline{\theta}$ (the state is high enough to make action A profitable), or if $\theta \geq \bar{\theta}$ (the state is high enough that the success of action A does not depend on player j 's choice).⁹ Thus, players face the following payoffs:¹⁰

	Success	Failure
Action A	$\theta - T$	$-T$
Action B	0	0

The state variable θ follows a normal distribution with mean μ_θ and variance σ_θ^2 . Players do not observe the realization of θ . However, each player $i = 1, 2$ observes a noisy private signal about it:

$$x_i = \theta + \sigma_i \varepsilon_i,$$

where $\sigma_i > 0$ and $\varepsilon_i \sim N(0, 1)$. The noise ε_i is *i.i.d.* across players and we denote by $\phi(\cdot)$ its probability density function, and $\Phi(\cdot)$ its cumulative distribution function. The precision of the signal that each player receives is determined by its standard deviation, σ_i .¹¹ In the model with exogenous information structures $\sigma_i = \sigma_j = \sigma$.

⁸The details of the model with endogenous information structures can be found in Section B of the appendix.

⁹ $\bar{\theta}$ and $\underline{\theta}$ define, respectively, upper and lower dominance regions for the fundamental.

¹⁰One can think of a number of applications where θ represents the relevant fundamentals. For example, it can represent the return to a risky investment or the gain from overthrowing an oppressive regime. In the first case, the action of players would be to invest (A) or not (B). In the second case, the action would be to engage in a political protest (A) or not (B). We can think of T as an investment cost in the first case, or as an opportunity cost in the second.

¹¹The precision of a random variable that follows normal distribution is defined as the inverse of its variance. In what follows we use higher precision, lower variance, or lower standard deviation interchangeably to describe the informativeness of signals.

2.2 Equilibrium

The equilibrium of the global game with heterogenous precisions follows the same intuition as the standard global games with homogenous precisions (see, for example, Morris and Shin (2003)). Since this intuition has been established in the literature, we skip the intermediate steps and refer an interested reader to the online appendix for details.¹²

Let σ_1, σ_2 be the standard deviations of the signals of player 1 and player 2, respectively. Suppose that each player uses a monotone strategy by which he chooses action A if his signal x_i is larger than a threshold x_i^* , and he chooses action B otherwise. Given player j 's threshold x_j^* , player i 's threshold x_i^* is determined by the solution to the following indifference condition:

$$E\left[\theta \Pr(x_j > x_j^* | \theta) \mid x_i^*, \theta \in [\underline{\theta}, \bar{\theta}]\right] + E\left[\theta \mid x_i^*, \theta > \bar{\theta}\right] - T = 0 \quad (1)$$

which states that at the threshold signal x_i^* player i is indifferent between taking action A and taking action B .

The equilibrium in monotone strategies is described by a pair of thresholds $\{x_1^*, x_2^*\}$ such that for each player $i = 1, 2$, the threshold x_i^* solves Equation (1) given that player $j \neq i$ follows a threshold x_j^* . Similar to the results in the literature, the coordination game has a unique equilibrium as long as the standard deviations of private signals are small enough compared to the standard deviation of the prior, σ_θ . Moreover, as the precision of the signals increases, the optimal thresholds converge to the risk-dominant threshold, which in our case is equal to $2T$.¹³ The next proposition formalizes these results.

Proposition 1 *Let $\{\sigma_1, \sigma_2\}$ be the standard deviations of players' signals. There exists a unique, dominance solvable equilibrium in which both players use threshold strategies characterized by $\{x_1^*(\sigma), x_2^*(\sigma)\}$ if either:*

1. $\frac{\sigma_i}{\sigma_\theta} < K_i(\underline{\theta}, \bar{\theta}, \mu_\theta)$, $i = 1, 2$ holds, for any pair of (σ_1, σ_2) ,¹⁴ or
2. $\sigma_\theta > \bar{\sigma}_\theta$, where $\bar{\sigma}_\theta$ is determined by the parameters of the model.

Moreover as $\sigma_i \rightarrow 0$, $\sigma_j \rightarrow 0$ and $\frac{\sigma_i}{\sigma_j} \rightarrow c \in \mathbb{R}$ this equilibrium converges to the risk-dominant equilibrium of the complete information game.

¹²The proofs of existence and uniqueness of equilibrium follow standard methods in the literature on global games. However, our results are not simple corollaries of the existing papers. Most of the literature deals with either the limiting case when the noise in the signals vanishes (Carlsson and van Damme, 1993, Frankel et al., 2003) or with a continuum of players (as in Szkup and Trevino, 2015). The fact that there are a discrete number of asymmetric players and that the prior is proper results in significantly more complex conditions for uniqueness than typically encountered in the literature.

¹³The risk-dominant threshold in our game is the optimal threshold when a player assigns equal probability to the other player taking either action.

¹⁴The derivation of the expression for $K_i(\underline{\theta}, \bar{\theta}, \mu_\theta)$ can be found in the online appendix.

The model with exogenous information structures corresponds to the case where $\sigma_i = \sigma_j = \sigma$, and is similar to Carlsson and Van Damme (1993) or Morris and Shin (2003).

We will test experimentally the predictions of Proposition 1 in an effort to understand whether subjects use threshold strategies and how these thresholds depend on the informativeness of signals (determined by σ). In particular, the treatment variations in the signal precision will allow us to approximate the path towards complete information with the objective to document empirically whether thresholds “converge” to a specific equilibrium of the complete information game. In the treatments with costly information acquisition we endogenize the choice of signal precision to study these questions when the path towards complete information is endogenously determined.

2.3 Costly information acquisition

The model with endogenous information structures is a two-stage model where players first acquire information and then play the coordination game described above. In the first stage each player decides how much information about θ to acquire by choosing the standard deviation of his signal, $\sigma_i \in (0, \underline{\sigma}]$. If a player chooses not to acquire information he will observe a signal with a default standard deviation $\underline{\sigma}$. The cost of choosing a standard deviation σ_i is $C(\sigma_i)$. The function $C(\cdot)$ is continuous, with $C(\underline{\sigma}) = 0$, $C'(\underline{\sigma}) = 0$, $C'(\sigma_i) < 0$, $C''(\sigma_i) > 0$, for all $\sigma \in (0, \underline{\sigma})$, and $\lim_{\sigma \rightarrow 0^+} C'(\sigma_i) = -\infty$. These assumptions imply that the cost of decreasing the standard deviation is increasing, convex, that infinitesimal information acquisition is costless, and that the marginal cost of acquiring better information converges to infinity as σ tends to 0. Once players have chosen the precision of their signal, they play the coordination game as described above.

The model with endogenous information is solved by backward induction. In the second stage, equilibrium behavior follows Proposition 1. In the first stage, players compare the benefit of being better informed to the cost of information. Let $B(\sigma_i; \{\sigma_j, \sigma'_i\})$ denote the benefit of choosing standard deviation σ_i to player i when player j chooses standard deviation σ_j . Player j expects player i to choose standard deviation σ'_i and both players behave optimally in the coordination stage given their beliefs (i.e., they follow the monotone strategies described above). The standard deviations $\{\sigma_1^*, \sigma_2^*\}$ constitute an equilibrium of the two-stage game if for each $i = 1, 2$ and each $j \neq i$

$$\frac{\partial}{\partial \sigma_i} B(\sigma_i^*; \{\sigma_j^*, \sigma_i^*\}) = C'(\sigma_i^*)$$

Section B of the Appendix contains a full solution and proof of existence of a symmetric equilibrium for the game with costly information acquisition.¹⁵

¹⁵In a monotone equilibrium the benefit of a higher precision comes from the reduction in the expected cost of two types of mistakes: taking action A when action A is unsuccessful or when $\theta < T$ (Type I mistake) and choosing action B when action A is successful and $\theta > T$ (Type II mistake). See Szkup and Trevino (2015) for an in-depth discussion of these mistakes in a global game with a continuum of players.

2.4 Strategic uncertainty

In this section we briefly describe how strategic uncertainty in this game depends on the signal precision, which is key to understand why equilibrium thresholds converge to the risk dominant equilibrium as the signal noise vanishes. We will revisit these concepts in Section 5 when we extend the model to reconcile the theory with our experimental findings.

One typically refers to strategic uncertainty as the uncertainty about the actions of other players (see Van Huyck et al., 1990, or Brandenburger, 1996). According to this definition, strategic uncertainty is high if a player is very uncertain about the behavior of others. In our model the key object that allows us to measure strategic uncertainty faced by player i is $\Pr(x_j > x_j^* | x_i)$, which represents the probability that player i assigns to player j taking action A . If $\Pr(x_j > x_j^* | x_i)$ is close to $1/2$ then player i deems each action by player j almost equally likely and hence faces high strategic uncertainty. On the other hand, if $\Pr(x_j > x_j^* | x_i)$ is close to 0 or 1 then he expects player j to take a particular action, thus he faces little strategic uncertainty.

The extent of strategic uncertainty that players face in the game varies with σ_i . As σ_i decreases, player i 's signal is closer to the state θ , so he is able to better estimate player j 's signal. Thus, if he receives a high (low) signal, he believes that player j also receives a high (low) signal and he assigns a higher probability to player j choosing action A (action B). Consider now the case where player i observes the signal $x_i = x_j^*$, that is, a signal equal to his opponent's threshold. In this case, an increase in the precision of x_i will increase strategic uncertainty since $\Pr(x_j > x_j^* | x_i = x_j^*)$ converges monotonically to $1/2$. Thus, for signals around x_j^* , the strategic uncertainty faced by player i increases as σ_i decreases. This leads to the limit result in Proposition 1, first shown by Carlsson and Van Damme (1993) for homogenous signal distributions.

To sum up, in this model a reduction of fundamental uncertainty (characterized by an increase in the precision of private information) increases strategic uncertainty for intermediate signals (i.e., signals in the neighborhood of x_i^* and x_j^*) and decreases strategic uncertainty for high or low signals. In fact, it is the increase in strategic uncertainty for intermediate signals that determines how player i adjusts his threshold when fundamental uncertainty decreases. Our experimental results will allow us to test whether this relation between strategic and fundamental uncertainty is consistent with the behavior of subjects.¹⁶

¹⁶Morris and Shin (2002, 2003) suggest using rank beliefs to measure strategic uncertainty, where rank beliefs are defined as the beliefs that player i attaches to his opponent receiving a higher or lower signal than him. Despite the theoretical advantages of studying rank beliefs (Morris, Shin and Yildiz, 2016), we believe that measuring strategic uncertainty using players' beliefs regarding the action of the other players is more intuitive in an experimental setting.

3 Experimental design

In this section we describe our experimental design and the predictions of the model that we test. We implement a between subjects design that allows us to directly compare the behavior of subjects across treatments.

There are three main dimensions in which our treatments vary: The nature of the information structure (exogenous and endogenous), the precision of the private signals (for exogenous information structures), and the way in which subjects choose actions (direct action choice and strategy method). Table 1 summarizes our experimental design.

When the information is given to subjects exogenously (first 7 treatments in Table 1) we vary the precision of the private signals in the following way: complete information (standard deviation of 0), high precision (standard deviation of 1), medium precision (standard deviation of 10), and low precision (standard deviation of 20). In the treatments with an endogenous information structure (treatments 8 and 9 in Table 1) subjects have to choose the precision of their private signals from the set of standard deviations of $\{1, 3, 6, 10, 16, 20\}$. In the treatments with direct action choice subjects choose an action (A or B , i.e., risky or safe) after observing their signal, as in the model described above. In the treatments with the strategy method for action choices subjects have to report a cutoff value such that they would choose the risky action (A) if their signal is higher than this cutoff and the safe action (B) if their signal is lower than the cutoff they report. Eliciting thresholds in this way allows us to observe the evolution of thresholds over time.

We also run treatments with exogenous information structures where we elicit first and second order beliefs (last 3 rows in Table 1).^{17,18}

¹⁷In each round, we elicit subjects' first order beliefs about the state θ after observing their signal. Once they have chosen an action, we elicit their second order beliefs about the probability they assign to their opponent having chosen actions A and B . Both elicitation of beliefs are incentivized using a quadratic scoring rule.

¹⁸Given our interest in studying fundamental and strategic uncertainty, one could think of eliciting beliefs of our subjects to learn their beliefs about the state and about the action taken by their opponent. However, we first follow a revealed preference approach and focus only on choice data to study how the behavior of the game varies with the information structure. We do not elicit beliefs in the treatments corresponding to the first 9 rows of Table 1 because we do not want to alter the individual reasoning of subjects by drawing attention to fundamental and strategic uncertainty. However, once we identify systematic deviations from the theory in our data, we hypothesize that biases in belief formation might be behind these results and test this hypothesis directly with these additional treatments.

TREATMENT	INFORMATION	SIGNAL PRECISION	CHOICE OF ACTION	BELIEF ELICITATION	# SUBJECTS
1. COMPLETE INFORMATION	EXOGENOUS	PERFECT (sd = 0)	DIRECT	NO	22
2. EXOGENOUS / HIGH / ACTION	EXOGENOUS	HIGH (sd = 1)	DIRECT	NO	38
3. EXOGENOUS / HIGH / STRATEGY	EXOGENOUS	HIGH (sd = 1)	STR. METHOD	NO	20
4. EXOGENOUS / MEDIUM / ACTION	EXOGENOUS	MEDIUM (sd = 10)	DIRECT	NO	40
5. EXOGENOUS / MEDIUM / STRATEGY	EXOGENOUS	MEDIUM (sd = 10)	STR. METHOD	NO	24
6. EXOGENOUS / LOW / ACTION	EXOGENOUS	LOW (sd = 20)	DIRECT	NO	44
7. EXOGENOUS / LOW / STRATEGY	EXOGENOUS	LOW (sd = 20)	STR. METHOD	NO	20
8. ENDOGENOUS / ACTION	ENDOGENOUS	sd ∈ {1,3,6,10,16,20}	DIRECT	NO	40
9. ENDOGENOUS / STRATEGY	ENDOGENOUS	sd ∈ {1,3,6,10,16,20}	STR. METHOD	NO	44
10. EXOGENOUS / HIGH / BELIEFS	EXOGENOUS	HIGH (sd = 1)	DIRECT	YES	22
11. EXOGENOUS / MEDIUM / BELIEFS	EXOGENOUS	MEDIUM (sd = 10)	DIRECT	YES	20
12. EXOGENOUS / LOW / BELIEFS	EXOGENOUS	LOW (sd = 20)	DIRECT	YES	16

Table 1: Experimental design

Each session of the experiment consists of 50 independent and identical rounds. The computer randomly selects five of the rounds played and subjects are paid the average of the payoffs obtained in those rounds, using the exchange rate of 3 tokens per 1 US dollar.¹⁹

Subjects are randomly matched in pairs at the beginning of the session and play with the same partner in all rounds.²⁰ To avoid framing effects the instructions use a neutral terminology. To avoid bankruptcies subjects enter each round with an endowment of 24 tokens. From Table 2 we can see that in the treatments with costly information acquisition even if subjects choose the precision with the highest cost the lowest payoff they can get in a round is 0, in case of miscoordination.

Before starting the first paying round subjects have access to a practice screen where they can generate signals (for the different available precisions if information is endogenous) and they are given an interactive explanation of the payoffs associated to each possible action, given their signal and the underlying state θ .

After each round subjects receive feedback about their own private signal, their choice of action, the realization of θ , how many people in their pair chose action A , whether A was successful or not, and their individual payoff for the round. In addition, in the treatments with endogenous information subjects observe precision choices and can access the history of precision choices made

¹⁹Note that while subjects' earnings are potentially unbounded, high payoffs are extremely unlikely to occur (e.g., the probability that a subject will earn more than \$60, assuming he always plays optimally, is about 0.1%).

²⁰We choose fixed pairs, as opposed to random pairs, to be able to study coordination of information choices over time. Due to the complexity of the setup with endogenous information structures, subjects might need time to learn because information choices are not observable and equilibrium assumes that beliefs about information choices are correct. Subjects choose consistently their level of precision, which "fixes" the beliefs about the information choices within a pair. For this reason we believe that fixed pairs are better suited to study this environment. In order to test if our results were due to the matching protocol, we run an additional session with random matching in each round and high signal precision and find no significant effect resulting from the matching procedure.

by both pair members over the previous rounds by pressing a button.

The experiment was conducted at the Center for Experimental Social Science at New York University using the usual computerized recruiting procedures. Each session lasted from 60 to 90 minutes and subjects earned on average \$30, including a \$5 show up fee. All subjects were undergraduate students from New York University.²¹ The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007). There were a total of 18 sessions and 350 participants.

Our experiment is related to the work of Heinemann et al. (2004) (HNO04 henceforth) who test the predictions of the model by Morris and Shin (1998) in the laboratory.²² It is also related to Cabrales, Nagel, and Armenter (2007) and Duffy and Ochs (2012).

3.1 Theoretical predictions for the experiment

The theoretical model is governed by a set of parameters $\Theta = \{\mu_\theta, \sigma_\theta, \underline{\theta}, \bar{\theta}, T, \{\sigma_i\}, \{C(\sigma_i)\}\}$. In the experiment:

- The state θ is randomly drawn from a normal distribution with mean $\mu_\theta = 50$ and standard deviation $\sigma_\theta = 50$.
- The coordination region is for values of $\theta \in [0, 100)$, that is $\underline{\theta} = 0$ and $\bar{\theta} = 100$.
- The cost of taking action A is $T = 18$.
- For the treatments with endogenous information structures, precision choices and the associated costs are presented in the form of a menu of 6 precision levels, standard deviations, and costs:²³

Precision level	1	2	3	4	5	6
Standard deviation	1	3	6	10	16	20
Information Cost	6	5	4	2	1.5	1

Table 2: Precision choices

²¹Instructions for all treatments can be found at http://econweb.ucsd.edu/~itrevino/pdfs/instructions_st.pdf.

²²Unlike HNO04, our focus is to understand how behavior in a global game depends on detailed exogenous and endogenous variations of the information structure. Our experiment also differs from HNO04 in terms of implementation. HNO04 use uniform distributions for the state and for private signals and they give subjects in each round a block of 10 independent situations (signals) and subjects have to choose an action for each signal before getting feedback. They then get feedback about the 10 choices and move on to the next round where they face a new block of 10 decisions. They have 16 rounds of 10 situations each. Additionally, each game of HNO04 consisted of 15 players, as opposed to our two-player case.

²³In the remainder of the paper, we will refer to information choices as precision choices to be consistent with the language used in the implementation of the experiment. We will use the term precision as a qualitative measure of informativeness of the signals, that is we will compare levels of precision (low, medium, high), and not magnitudes of standard deviations.

We decided not to have a default precision chosen for subjects in order to avoid status quo biases. The reason to introduce a discrete choice set for precisions was to simplify the choice for subjects and the data analysis. We believe six is a reasonable number of options to observe dynamics in the level of informativeness that subjects choose, without losing statistical power.

Given these parametric assumptions we characterize the predictions of the model in the form of two main hypotheses to be tested with our experiment:

Hypothesis 1 (Exogenous information)

- a) *Subjects use equilibrium threshold strategies in treatments with exogenous information structures and incomplete information.*
- b) *Under complete information, subjects behave in accordance to the theoretical prediction of multiplicity of equilibria.*
- c) *Thresholds are increasing in precision and tend towards the risk-dominant threshold.*

Hypothesis 2 (Endogenous information)

- a) *Subjects use threshold strategies for their preferred precision choices.*
- b) *Subjects choose the unique equilibrium precision and threshold.*
- c) *Thresholds are increasing in precision choices and tend towards the risk-dominant threshold.²⁴*

Hypothesis 1 pertains to the treatments where subjects are exogenously endowed with the same signal precision or where they perfectly observe the state (treatments 1-7). Given the parameters used in the experiment, the equilibrium threshold when subjects observe signals with high precision (standard deviation of 1) is 35.31, for medium precision (sd of 10) it is 28.31, and with a low precision (sd of 20) it is 18.73. When subjects have complete information about θ , the theory suggests multiple equilibria.

Hypothesis 2 pertains to the treatments with endogenous information structures via costly information acquisition. Part (a) of the hypothesis states that subjects choose a unique equilibrium in threshold strategies, for a given precision choice. Part (b) aims to test the unique symmetric equilibrium prediction for the parameters used in the experiment, which corresponds to coordinating on choosing precision level 4 (sd of 10) and setting a symmetric threshold at 28.31. Implicit in this prediction is that precision choices are strategic complements, which leads players to coordinate on both precisions and actions. Given the parameters in the experiment, in part (c) of Hypothesis 2 we test the comparative statics of the thresholds in the coordination game with respect to precision choices in the first stage, in case subjects do not choose the equilibrium precision.

²⁴In general, in the model thresholds are increasing in the precision of information if μ_θ is high relative to T and decreasing when μ_θ is low relative to T . Given our choice of parameters, we are in the former case.

To understand the intuition behind the predictions for comparative statics of thresholds and precisions, note that thresholds will be low in general when μ_θ is high relative to T since this makes the risky action more likely to succeed in expectation.²⁵ This is stronger when the precision of signals is low, since in this case players assign a high weight to the prior. Thus, for a low precision of signals the model predicts low thresholds for our choice of parameters. An increase in the precision of signals has two effects on players' behavior. First, they assign a lower weight to the information contained in the prior. Second, it increases the correlation between the players' signals, making it harder for a player to predict whether his opponent observes a signal higher or lower than his own, which leads to higher strategic uncertainty for intermediate signals (close to the thresholds). For our choice of parameters, this means that thresholds will increase as the precision of signals increases. Finally, in the limit, as signals become perfectly informative, this increase in strategic uncertainty leads players to choose the risk dominant equilibrium of the game, as in the limit they assign probability 1/2 to the other player observing a signal higher or lower than their own.

4 Experimental results

In this section we first explain our methodology to estimate thresholds, since they are the relevant objects to study our hypotheses. We then investigate each of our hypotheses, and then we summarize the departures from the theoretical predictions found in the data to facilitate the introduction of the extended model in Section 5.²⁶

4.1 Estimation of thresholds

We say that a subject's behavior is consistent with the use of threshold strategies if the subject uses either perfect or almost perfect thresholds. A perfect threshold consists of taking the safe action B for low values of the signal and the risky action A for high values of the signal, with exactly one switching point (the set of signals for which a subject chooses A and the set of signals for which he chooses B are disjoint). For almost perfect thresholds, subjects choose action B for low signal values and action A for high signal values, but we allow these two sets to overlap for at most three observations. These two types of behavior are illustrated in Figure 5 in the appendix.

Once we identify the subjects who use threshold strategies, we use two different methods to

²⁵See Heinemann and Illing (2002), Bannier and Heinemann (2005), Iachan and Nenov (2014), or Szkup and Trevino (2015) for detailed theoretical discussions of comparative statics in global games with respect to precision of information.

²⁶Most of the results reported from the experiment pertain to the last 25 rounds to allow for behavior to stabilize, unless otherwise specified. This is particularly important for the treatments with endogenous information since subjects first have to stabilize in their precision choices before we can estimate thresholds.

estimate their thresholds. For the first method we pool the data of all the subjects who use thresholds in each treatment and fit a logistic function with random effects (RE) to determine the probability of taking the risky action as a function of the observed signal.²⁷ The cumulative logistic distribution function is defined as

$$\Pr(A) = \frac{1}{1 + \exp(\alpha + \beta x_i)}$$

We estimate the mean threshold of the group by finding the value of the signal for which subjects are indifferent between taking actions A and B , that is, the value of the signal for which they would take the risky action with probability $\frac{1}{2}$, which is given by $-\frac{\alpha}{\beta}$. As pointed out by HNO04, the standard deviation of the estimated threshold, $\frac{\pi}{\beta\sqrt{3}}$, is a measure of coordination and reflects variations within the group. We call this the Logit (RE) method.

For the second method we take the average, individual by individual, between the highest signal for which a subject chooses the safe action and the lowest signal for which he chooses the risky action. This number approximates the value of the signal for which he switches from taking one action to taking the other action. Once we have estimated individual thresholds this way, we take the mean and standard deviation of the thresholds in the group. We refer to this estimate as the Mean Estimated Threshold (MET) of the group.

4.2 Hypothesis 1: Exogenous information

To study part (a) of Hypothesis 1, we use the procedure described above to determine how many subjects use thresholds strategies. We find that the behavior of 93.44% of subjects is consistent with the use of threshold strategies.²⁸ Table 3 shows the mean estimated thresholds for the different treatments with exogenous information structures. Standard deviations are reported in parenthesis. While we find that the vast majority of subjects uses threshold strategies, these thresholds differ substantially from the thresholds predicted by the theory. In particular, we reject to the 1% and 5% levels of significance that the thresholds estimated for the high precision treatment coincide with the theoretical equilibrium threshold of 35.31, using the MET and logit methods, respectively. For the treatments with medium and low precision, we reject to the 1% level that subjects play the equilibrium thresholds of 28.31 and 18.73, respectively, using both methods.²⁹

²⁷For the treatment with endogenous information structures we pool the data of subjects according to the level of precision chosen.

²⁸In particular, 97.37% of subjects use thresholds for high precision, 92.5% for medium precision, and 90.91% for low precision. This result is qualitatively similar to HNO04, even if HNO04 use a different metric to measure the use of threshold strategies and have 10 decision situations in each round of the experiment.

²⁹The thresholds estimated for high precisions are different to the risk dominant threshold to the 1% level. For the treatments with medium and low precision we cannot reject the hypothesis that the estimated thresholds coincide with the risk dominant threshold of 36.

	Complete info	High precision	Medium precision	Low precision
Logit (RE)	22.01 (7.15)	27.61 (5.86)	40.16 (9.13)	35.79 (9.00)
MET	21.07 (11.85)	27.42 (19.16)	40.37 (18.77)	36.23 (23.36)
Equilibrium x^*		35.31	28.31	18.73
Risk dominant eq.		36	36	36

Table 3: Estimated thresholds and equilibrium predictions, exogenous information

We also reject part (b) of Hypothesis 1 with our data. In the treatment with complete information all of our subjects use threshold strategies, and thus behave in accordance with a unique equilibrium. This is consistent with the findings of HNO04, who document that subjects play threshold strategies in coordination games with complete information. Moreover, the mean estimated threshold for this treatment is close to the efficient threshold of 18 (with 77.27% of subjects in the complete information treatment using exactly the efficient threshold). This is also consistent with Charness, Feri, Melendez-Jimenez, and Sutter (2014) who show efficient play in coordination games under complete information.

Finally, we also reject part (c) of Hypothesis 1. The thresholds of subjects that are given either a medium or a low precision exogenously are not statistically different from each other, but they are statistically higher than the thresholds under high precision. This means that the estimated thresholds are non-increasing on the precision of information and seem to display a convergence towards the efficient threshold obtained under complete information, and not the risk dominant, as the theory predicts.³⁰ That is, a high precision of signals leads to more successful coordination than what is suggested by the theory.³¹ As is shown in Table 9 in the appendix, this deviation is payoff improving.

4.3 Hypothesis 2: Endogenous information

In order to investigate part (a) of Hypothesis 2 we need to first establish stability of individual precision choices when information is endogenously chosen by subjects. We find that subjects choose, on average, the same precision for more than 22 out of the last 25 rounds of the experiment. Figure 6 in the appendix shows the transition matrix of precision choices in the last 25 rounds of

³⁰We abuse language slightly and talk about convergence in behavior when the signal noise decreases. Given the discrete nature of experiments, it is not correct to make statements about convergence. We are aware of this limitation, but we use this term to give an intuitive interpretation to our results.

³¹Similar efficient thresholds are found in a robustness check with high precision and random rematching in every round, suggesting that the efficient play that we observe in the experiment when signals are very precise is not due to repeated game effects.

the experiment.³² The entry a_{ij} of the matrix shows the probability of choosing precision level j in $t + 1$, given that a subject chose precision level i at t , for $i, j \in [1, 6]$ and $t > 25$. By looking at the diagonal entries of the transition matrix, we can see that most precision levels (except for level 5) are absorbent states.³³ Given this stability result, we characterize subjects by their preferred precision choice.

Table 4 shows the percentage of subjects that choose each precision for the last 25 rounds of the experiment.³⁴ Notice that the most popular precision choice is the equilibrium precision (level 4).

Precision level	Standard deviation	Cost	Precision choices in last 25 rounds
1	1	6	14.7%
2	3	5	3.7%
3	6	4	18.4%
4*	10	2	36.9%
5	16	1.5	3.9%
6	20	1	22.4%

Table 4: Precision choices in the last 25 rounds, endogenous information

To test part (a) of Hypothesis 2, we analyze the use of threshold strategies, given each subject's preferred precision choice. We find that 100% of subjects choosing precision levels 1, 2, or 3 use threshold strategies. For precision level 4, 93.75% of the subjects use threshold strategies. For precision level 6, 75% of the subjects use threshold strategies and 25% always choose either action A or B . This suggests that when subjects invest in more precise information their behavior is more likely to be consistent with the use of threshold strategies. In total, 90% of the subjects in the endogenous information treatment with direct action choice use threshold strategies for their most preferred precision choice. This implies that the theoretical prediction of subjects using threshold strategies is robust also under endogenous information.

To study part (b) of Hypothesis 2, we first look at how individual decisions in the game depend on individual precision choices. Then, we study behavior at the pair level to estimate thresholds, since they depend on both precision choices within a pair.

³²This includes precision choices in treatments with a direct action choice and with the strategy method. We aggregate the data because the distributions of precision choices are not statistically different between these two treatments. This was expected since the treatment effect is in the second stage of the game.

³³Less than 5% of subjects chose precision 5 and their behavior in the second stage was mostly random.

³⁴Precision choices in the first rounds are not very dissimilar to the precision choices portrayed in Table 4. In particular, if we compare the choices of the first 5 rounds with the choices of the last 5 rounds of the experiment, we observe the following proportion of choices, by precision level (the first number corresponds to the first 5 rounds and the second to the last 5 rounds). Level 1: 16.2% vs 13.8%; level 2: 11.9% vs 4%; level 3: 25% vs 19%; level 4: 21.4% vs 36.8%; level 5: 4.8% vs 4%; level 6: 20.7% vs 22.4%. We observe the highest shift in precision choices to be in favor of the equilibrium precision level 4.

Figure 1 plots the cumulative distribution function (pooled over all subjects) to illustrate the probability of choosing the risky action for each signal realization, by precision levels. The value of the signal for which subjects choose the risky action with probability 0.5 determines a threshold. Looking at the intersection of the curves corresponding to the different precision levels with the 0.5 horizontal line, from left to right, we can see that, in general, thresholds are larger for lower precisions. This suggests that the subjects who acquire more precise information choose the risky action more often in an effort to coordinate, which is consistent with our findings under exogenous information, but in stark contrast to the theoretical predictions. We also see that lower precision levels exhibit less steep CDFs, indicating higher dispersion among the subjects that choose lower precisions.

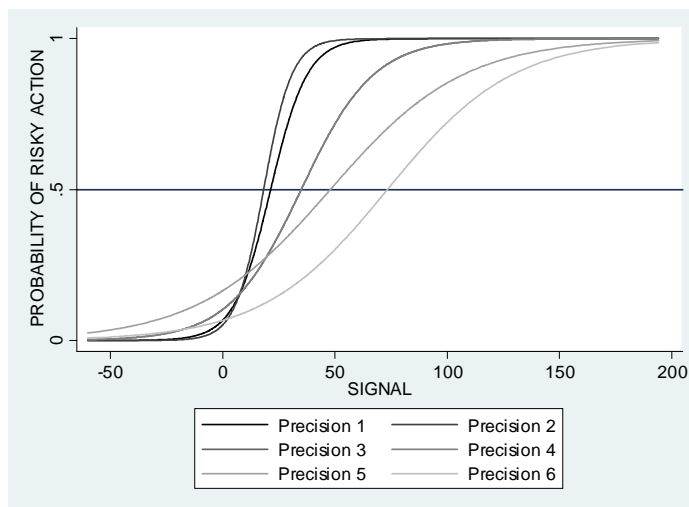


Figure 1: Probability of taking the risky action, by precision choices

We perform two regression estimations for the treatments with endogenous information (direct action choice and strategy method, respectively) and find strong support for the finding that as subjects choose higher precisions they try to coordinate more often on the risky action (see Tables 10 and 11 in the appendix).

In order to compare thresholds to equilibrium predictions we need to categorize precision choices within a pair, since thresholds depend on the precision choices of both pair members. We define individual convergence in precision as a situation where a subject chooses the same precision level for the last 25 rounds, with at most three deviations. We say that a pair exhibits non-stable behavior if at least one of its members does not converge individually in his precision choice. A pair that has stability but not convergence is a pair in which both members converge individually in their own precision choices, but the levels at which they converge are more than one level apart. We define weak convergence as pairs in which both members converge individually to a level of

precision and these two precision levels are at most one level away from each other. We say that a pair exhibits full convergence if both members converge individually to the same level of precision for the last 25 rounds of play.³⁵

To estimate thresholds we focus on weak and full convergence and we restrict our attention to pairs that coordinate on high precision (levels 1 and 2), medium precision (levels 3 and 4), and low precision (levels 5 and 6). These correspond to the diagonal entries of Table 5, which summarizes the combinations of precision choices across pairs that exhibit weak convergence. Approximately two thirds of the total number of pairs exhibit weak convergence in precision, with the majority converging to a medium precision, which corresponds to the theoretical prediction.

	High	Medium	Low
H	10.00%	13.81%	3.05%
M		40.00%	16.67%
L			16.48%

Table 5: Weak convergence of precision choices, endogenous information

With these results in hand, in Table 6 we compare the mean thresholds of pairs that converge to high, medium, and low precision to the thresholds predicted by the theory.³⁶ As we can see, subjects who choose the equilibrium precision (medium) choose, on average, the equilibrium threshold. In particular, we cannot reject the hypothesis that the thresholds estimated for these subjects using the MET and logit methods are different from the equilibrium threshold of 28.31. This means that the subjects that coordinate on a medium precision behave on average in accordance to the unique equilibrium suggested by the theory, unlike those who converge to either a high or low precision. This supports part (b) of Hypothesis 2.

³⁵Table 12 in the appendix shows all the combinations of precision choices made by the different pairs in our experiment (for both treatments with endogenous information). The diagonal entries correspond to the pairs that exhibit full convergence.

³⁶Since we define weak convergence to high precision as pairs that converge to precision levels 1 or 2, medium precision as pairs that converge to precision levels 3 or 4, and low precision as pairs that converge to precision levels 5 or 6, for each precision (high, medium, or low) we include the two predictions that correspond to each of the precision levels (1 and 2, 3 and 4, or 5 and 6), as well as the risk dominant equilibrium, i.e. the threshold prediction when the signal noise converges to zero.

	Complete info	High precision		Medium precision		Low precision	
RE Logit	22.01 (7.15)	24.99 (7.44)		30.30 (12.32)		74.62 (21.48)	
MET	21.07 (11.85)	25.29 (9.27)		27.84 (17.65)		50.65 (28.65)	
Equilibrium x^*		Info 1	Info 2	Info 3	Info 4	Info 5	Info 6
		35.31	33.88	31.61	28.31	22.82	18.73
Risk dominant eq.		36		36		36	

Table 6: Estimated thresholds and equilibrium predictions, endogenous information

However, we can see from Table 6 that part (c) of Hypothesis 2 is rejected. When the precision of information is endogenous the thresholds are actually decreasing in precision and, again, tend towards efficiency and not risk dominance. As shown in Table 13 in the appendix, the deviation that leads subjects to behave more efficiently under a higher signal precision is welfare improving, even if subjects pay a higher cost to acquire the most precise information.

To understand why the qualitative results with an endogenous information structure are starker than with an exogenous one, we turn our attention to the treatments where thresholds are elicited in every round using the strategy method. By looking at the evolution of thresholds over time (Figures 7 and 8 in the appendix) we see that the stability of individual thresholds and the convergence of thresholds within a pair depend on the level of precision only when subjects choose it. When precision is exogenously determined we see no significant differences in stability and convergence of thresholds across levels of precision. This suggests that we cannot exogenously manipulate subjects to have more stable individual behavior and better coordination within a pair by endowing them with more precise information. However, when the precision is endogenously chosen the subjects from pairs that choose a high precision show, on average, very stable individual thresholds and they converge to very similar thresholds within a pair. The individual and pair-wise convergence is weaker for subjects in pairs that choose medium precisions, and even more so for those who choose low precisions. We can interpret these results as reflecting the possibility that subjects “self-select” when allowed to choose the precision of their information, thus leading to starker departures from the theoretical predictions.

4.4 Summary of departures from the theoretical predictions

Our experimental analysis above highlights two main empirical departures from the theoretical predictions: (1) thresholds tend to decrease as signals become more precise, rather than increase as predicted by the theory, and (2) as the signal precision increases, thresholds tend towards the efficient threshold, rather than the risk dominant one. This systematic path to convergence towards efficiency illustrates an underlying force in the game that is not captured by the theory. These two

stark discrepancies in our data with respect to the theoretical predictions are illustrated in Figure 2, which plots the estimated thresholds for exogenous and endogenous information structures (solid and dashed blue lines, respectively) and the theoretical predictions for the different noise levels (red dash-dotted line) and a horizontal line at the risk dominant equilibrium.

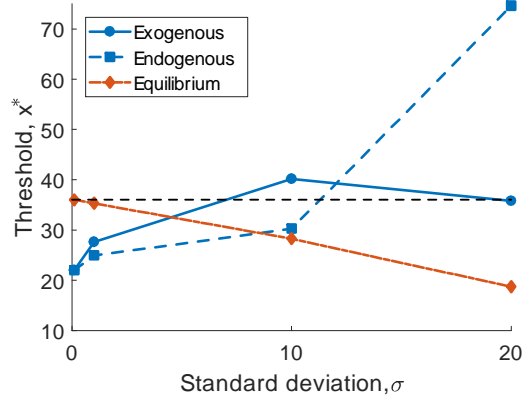


Figure 2: Theoretical predictions and estimated thresholds for exogenous and endogenous information structures.

It is worth stressing that these departures from the theoretical predictions are robust in the sense that we observe them across treatments with different information structures (exogenous and endogenous), different ways to estimate thresholds, and, as we show below, to belief elicitation (see Section 5.3.1).³⁷

5 Understanding departures from the theory

We have documented two empirical departures from the theory. In order to understand what could be driving these deviations we go back to the benchmark model of Section 2 and focus on what happens in the limit, as the signal noise vanishes, $\sigma_i \rightarrow 0$ for $i = 1, 2$. In this case, one can show that the indifference condition of player $i \in \{1, 2\}$ is approximately given by³⁸

$$\Pr(x_j \geq x_j^* | x_i^*) \mathbb{E}[\theta | x_i^*] = T \quad (2)$$

Equation (2) tells us that when players' signals are very precise their equilibrium behavior is determined by the fundamental uncertainty (captured by $\mathbb{E}[\theta | x_i^*]$) and the strategic uncertainty (captured by $\Pr(x_j \geq x_j^* | x_i^*)$) that players face in equilibrium. This suggests that, under the

³⁷The same patterns, albeit more noisy, appear when we restrict our attention to earlier rounds. This suggests that learning about the coordination motives of opponents is not the main driver behind these results.

³⁸One can show that for a sufficiently small σ , the *LHS* of Equation (1) is approximately equal $E[\theta | x_i^*] \Pr(x_j \geq x_j^* | x_i^*)$. This can be deduced from the fact that in the limit as $\sigma \rightarrow 0$, the indifference condition becomes $x_i^* \Pr(x_j \geq x_j^* | x_i^*) = T$.

assumption that players best-respond to their beliefs, deviations from the theoretical predictions can be driven by subjective (or biased) perceptions of fundamental and/or strategic uncertainty. We refer to these subjective perceptions as sentiments (see Angeletos and La’O, 2013; or Izmalkov and Yildiz, 2010). Note, however, that as $\sigma_i \rightarrow 0$ then $\mathbb{E}[\theta|x_i] \rightarrow x_i$, and thus the task of computing the expected value of θ becomes rather simple. This suggests that, at least in the case of high precision, departures from the theory observed in the experiment are unlikely to be driven by a biased perception of fundamental uncertainty due to sentiments, but rather by sentiments that affect the perception of strategic uncertainty. Motivated by this observation, we formulate the following hypothesis.

Hypothesis 3 *The behavior of subjects is driven by a biased perception of strategic uncertainty. Moreover, this bias leads subjects to perceive less strategic uncertainty when information is precise and more when information is imprecise.*

This hypothesis posits that the departures from the theory are driven by biases in the perception of strategic uncertainty that are determined by the precision of private information in the following way. Subjects are optimistic about other subjects taking the risky action to coordinate when information is very precise, and they are pessimistic when information is imprecise. In the next section we formalize this hypothesis and provide a particular justification for it in terms of biased second-order beliefs. We extend our baseline model to allow for sentiments that bias the beliefs about the action of the other player.

We then test Hypothesis 3 in two ways. We first ask whether the extended model can be used to rationalize the subjects’ behavior presented above. We then test the belief-based mechanism of Hypothesis 3 directly by running new experiments where we elicit first and second-order beliefs.³⁹

5.1 Model with sentiments

In this section, we extend the baseline model to accommodate a biased perception of strategic uncertainty. For simplicity, we focus on the case with exogenous information structures where all signals are drawn from the same distribution with mean θ and standard deviation σ . For each $i = 1, 2$ and $j \neq i$, let $\Pr_i(x_j \geq x_j^*|x_i)$ denote the subjective probability that player i assigns to player j taking the risky action. We consider a specific form of additive bias, or sentiment, where

³⁹In Section 5.4 below, and in Section B of the Appendix we discuss the suitability of alternative equilibrium notions that assume bounded rationality (such as QRE, limited depth of reasoning, correlated equilibria, or analogy-based equilibria) and argue that, unlike our extended model, these alternative models cannot fully organize our experimental findings.

$\Pr_i(x_j \geq x_j^* | x_i)$ is given by

$$\Pr_i(x_j \geq x_j^* | x_i) = \int_{\theta=-\infty}^{\infty} \left[1 - \Phi\left(\frac{x_j^* - \theta}{\sigma} - \alpha_k\right) \right] \phi(\theta | x_i) d\theta. \quad (3)$$

That is, α_k denotes player i 's bias when estimating the probability that the other player takes the risky action.⁴⁰ Let $\alpha_k \in \{\alpha_1, \alpha_2, \dots, \alpha_N\}$ where, $\alpha_k < \alpha_{k+1}$. Note that when $\alpha_k = 0$, player i forms correct (Bayesian) beliefs about player j 's choice of action. If $\alpha_k > 0$ then the probability that player i assigns to player j acting is biased towards taking the risky action (positive sentiments about coordination), while if $\alpha_k < 0$ the bias is towards player j taking the safe action (negative sentiments about coordination).

The sentiments captured by α_k 's can arise from biased second-order beliefs if, for example, player i believes that player j extracts the information from his private signal x_j in a biased way. That is, if player i believes that player j computes his posterior mean as if his signal was $x_j = \theta + \sigma\varepsilon_j + \sigma\alpha_k$.⁴¹ Thus, one can interpret α_k as measuring player i 's overconfidence (or underconfidence) about player j taking the risky action, which results from player i 's beliefs about player j 's overestimation (or underestimation) of the value of the state θ . Note that such an interpretation is consistent with the second part of Hypothesis 3 if players believe that on average $\alpha_k > 0$ when information is precise and $\alpha_k < 0$ otherwise.⁴²

We denote the types of players i and j as the tuples $\{\{x_i, \alpha_k\}, \{x_j, \alpha_l\}\}$ and assume that the presence of the opponent's bias is common knowledge but not its magnitude. Instead, player i believes that $\alpha_l \in \{\alpha_1, \alpha_2, \dots, \alpha_N\}$ and assigns probability $g(\alpha_l)$ to player j 's bias being equal to α_l , for each $l = 1, \dots, N$. Thus, each player is uncertain not only about the signal that the other player observes but also about the magnitude of his opponent's bias. Finally, while player i takes into account the possibility that player j has positive or negative sentiments about the probability that he himself takes the risky action, player i thinks that his own assessment of the probability that player j 's takes the risky action is objective.⁴³

Let $x_i^*(\alpha_k)$ denote the threshold above which player i takes the risky action when his bias is equal to α_k .

⁴⁰To economize on notation we denote the conditional distribution of θ given x_i by $\phi(\theta | x_i)$.

⁴¹We still refer to α_k as player i 's bias (and not player j 's) because the bias α_k is part of player i 's belief about player j 's beliefs.

⁴²The formulation of the subjective probability of Equation (3) is also consistent with other interpretations. For example, it could be the result of introducing uncertainty about the threshold used by player j in equilibrium. If we define the set of N possible thresholds that player j can use as $\{x_{jk}^*\}_{k=1}^N$, where $x_{jk}^* = x_j^* - \sigma\alpha_k$, then Equation (3) would correspond to the probability that player i ascribes to player j taking the risky action when he believes that player j 's equilibrium threshold is x_{jk}^* .

⁴³This is similar to the literature on overconfidence (see for example García et al., 2007).

Definition 2 (Sentiments equilibrium) *A pure strategy symmetric Bayesian Nash Equilibrium in monotone strategies is a set of thresholds $\{x_i^*(a_k)\}_{k=1}^N$ for each player $i = 1, 2$ such that for each i and each $k = 1, \dots, N$, the threshold $x_i^*(a_k)$ is the solution to*

$$\sum_{l=1}^N g(\alpha_l) \int_{\underline{\theta}}^{\bar{\theta}} \theta \left[1 - \Phi \left(\frac{x_j^*(\alpha_l) - \theta}{\sigma} - \alpha_k \right) \right] \phi(\theta | x_i^*(\alpha_k)) d\theta + \int_{\bar{\theta}}^{\infty} \theta \phi(\theta | x_i^*(\alpha_k)) d\theta = T \quad (4)$$

where $x_j^*(\alpha_l)$ is the threshold of player $j \neq i$ when he exhibits sentiments derived from the bias α_l , $l \in \{1, \dots, N\}$.⁴⁴

Notice that we now have $2N$ equilibrium conditions. Any set of thresholds $\{x_i^*(\alpha_k), x_j^*(\alpha_l)\}_{k,l=1,\dots,N}$ that solves this system of equations constitutes an equilibrium. The next proposition establishes that the extended model has a unique equilibrium with sentiments when the noise of the private signals is small enough.

Proposition 3 *Consider the extended model.*

1. *There exists $\bar{\sigma} > 0$ such that for all $\sigma \in (0, \bar{\sigma}]$ the extended model has a unique equilibrium in monotone strategies which is symmetric. In this equilibrium $x^*(\alpha_k) > x^*(\alpha_l)$ for all $k < l$ where $k, l \in \{1, \dots, N\}$.*
2. *As $\sigma \rightarrow 0$ we have*

$$x^*(\alpha_k) \rightarrow x^* \text{ for all } k \in \{1, \dots, N\} \quad (5)$$

where

$$x^* = \frac{T}{\sum_{k=1}^N g(\alpha_k) \Phi\left(\frac{\alpha_k}{\sqrt{2}}\right)}.^{45}$$

The above proposition states that the extended model has a unique equilibrium in monotone strategies when private signals are sufficiently precise, similarly to the standard model. In the unique equilibrium of the extended model the thresholds are decreasing functions of players' biases. This is intuitive since a higher bias α_k for player i reflects positive sentiments implying that, for any signal x_i , player i assigns a higher probability to player j taking the risky action. The second part of the proposition states that in the limit players use the same threshold regardless of the

⁴⁴Since player i has to take into account that the threshold of player j depends on the magnitude of his own bias, the indifference condition of player i with bias α_k includes the summation over all possible biases that player j may have.

⁴⁵More precisely, the above limit applies if and only if $T / (\sum_{k=1}^N g(\alpha_k) \Phi(\alpha_k / \sqrt{2})) \leq \bar{\theta}$. If $T / (\sum_{k=1}^N g(\alpha_k) \Phi(\alpha_k / \sqrt{2})) > \bar{\theta}$ then in the limit $x^*(\alpha_k) = \bar{\theta}$ for all $k = 1, \dots, N$. For more details see Section C of the Appendix.

magnitude of their bias. Thus, the dispersion of the thresholds of players with different sentiments decreases as the noise in the private signals disappears.⁴⁶

Different from the standard model, the threshold in the limit differs in general from the risk-dominant threshold ($2T$), depending on the distribution of the biases. For example if α_k 's increase when precision increases (sentiments lead to more optimism as information becomes more precise), the limiting threshold tends to the welfare efficient threshold T , while if α_k 's decrease as precision increases (sentiments lead to more pessimism as information becomes more precise) the limiting threshold tends to $\bar{\theta}$. If $\alpha_k = 0$ for all $k = 1, \dots, N$ then we recover the risk-dominant threshold in the limit, as in the standard model.

5.2 The extended model and the data

We now use our extended model to interpret the experimental results of Section 4. In particular, we compute the average value of α_k 's that would be consistent with the observed thresholds for the different precisions in the treatments with exogenous and endogenous information structures. The resulting α_k 's are depicted below.

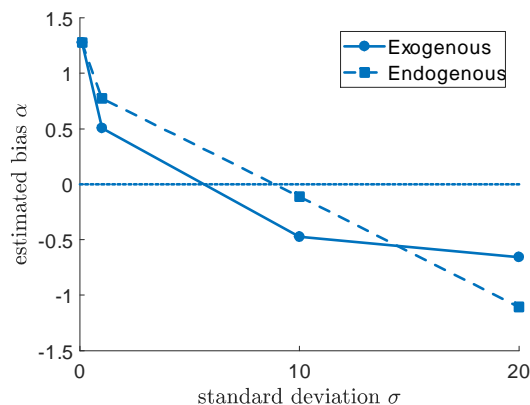


Figure 3: Average additive biases consistent with estimated thresholds.

Figure 3 shows that the low thresholds associated with a high precision (low standard deviation) in our data are consistent with a large positive bias, while the high thresholds associated with low precisions (high standard deviation) are consistent with a large negative bias. These results support Hypothesis 3.

⁴⁶To understand this, note that a player with a higher bias always uses a lower threshold. Thus, if thresholds were not converging to the same limit as $\sigma \rightarrow 0$ then a player with a higher bias who receives the threshold signal would expect both a lower payoff from the risky action when the risky action is successful and a lower probability that the risky action is successful compared to a player with lower bias who receives his threshold signal. This follows from the observation that as $\sigma \rightarrow 0$ each player expects that the other player observes the same signal. But this means that the expected payoff from taking risky action is lower for a player with high bias at his threshold signal than that of a player with low bias at his threshold signal, which is a contradiction.

However, even if Figure 3 suggests that the extended model can organize our experimental findings, we still have to test whether the mechanism based on sentiments is driving the results. Therefore, in what follows, we test Hypothesis 3 directly.

5.3 Experimental test of the extended model

To test Hypothesis 3 experimentally, we run additional sessions where we elicit subjects' first- and second-order beliefs. The experimental protocol of these sessions is identical to the one with exogenous information structures (for high, medium, and low precisions), except for two additional questions that asked subjects to report their best guess about the state θ (after observing their private signals) and the probability they assign to their opponent taking the risky action (after choosing their own action and before getting feedback).⁴⁷

The elicited beliefs allow us to understand the subjects' perception of fundamental uncertainty (via their reports about the value of the state) and of strategic uncertainty (via their reports about the probability of their opponent taking the risky action) and thus identify the type of sentiments that might arise under different signal precisions. We focus first on subjects' first-order beliefs and confirm our hypothesis that the observed departures from the theory are not driven by sentiments about fundamentals. We then analyze the elicited second-order beliefs and investigate whether subjects exhibited pessimism about the action of others when the signal precision is low and optimism when the signal precision is high.

5.3.1 Preliminaries

Before we analyze the elicited beliefs we estimate the thresholds for these additional sessions to confirm that the departures from the theory presented in Section 4 are also present in the new data set. As we can see in Table 7, when we elicit beliefs the estimated thresholds are decreasing in precision, consistent with our previous results. These results are starker than those of the treatments with exogenous precisions and no belief elicitation, and they are similar to the thresholds with endogenous information structures. Figure 9 in the appendix plots the thresholds under these three treatment conditions. These starker thresholds can be due to the salience created by the elicitation of beliefs. It has been documented that eliciting beliefs might affect the way subjects play the game by accelerating best-response behavior (see Croson, 2000, Gächter and Renner, 2010, or Rutström and Wilcox, 2009).⁴⁸ This is intuitive because belief elicitation forces subjects to think about fundamental and strategic uncertainty, thus putting more structure to their thought process. The thresholds of Table 7 will serve as the reference thresholds for the analysis that follows.

⁴⁷The elicitation of first and second order beliefs was incentivized using a quadratic scoring rule.

⁴⁸For a survey on belief elicitation, see Schotter and Trevino (2015).

	High precision	Medium precision	Low precision
Logit (RE)	19.24 (4.71)	28.76 (8.76)	40.72 (11.02)
MET	19.24 (7.64)	29.36 (22.61)	39.85 (27.92)
Equilibrium x^*	35.31	28.31	18.73
Risk dominant eq.	36	36	36

Table 7: Estimated thresholds and equilibrium predictions, belief elicitation.

5.3.2 Perception of fundamental uncertainty

Let θ_i^B denote the stated belief of subject i about θ . Table 8 shows the mean and median differences in absolute value between the stated beliefs about the state and the true state ($|\theta_i^B - \theta|$), and between the stated beliefs about the state and the observed signal ($|\theta_i^B - x_i|$), for each level of precision. As expected, subjects rely on their signals more and form more accurate beliefs about the state when signals are very precise, and this decreases as signals become noisier.

	$ \theta_i^B - \theta $			$ \theta_i^B - x_i $		
	High	Medium	Low	High	Medium	Low
Mean	1.19	10.41	19.59	0.83	7.02	12.24
Median	0.88	7.38	17.45	0.55	4.81	9.99
St. dev.	3.79	28.33	14.57	3.84	27.5	12.72

Table 8: Differences in absolute value between first order beliefs, the true state, and observed signals, by precision.

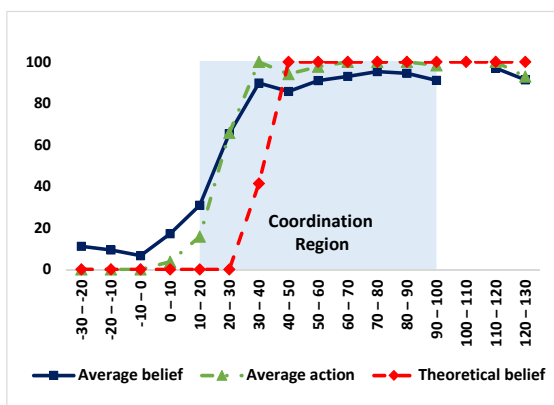
To evaluate the accuracy of subjects' first order beliefs, we estimate the weights assigned to the private signal and to the prior when forming beliefs about θ . We can see in Table 14 in the appendix that subjects put higher weights on the private signal as its precision increases and that the mean and median weights are not statistically different to the Bayesian weights for high and medium precision. For a low precision subjects seem to assign a higher weight to private signals than Bayes' rule implies.⁴⁹

To summarize, the fact that subjects set significantly lower thresholds than those suggested by equilibrium when the signal precision is high is unlikely to be explained by biases in beliefs about fundamentals, since in this case subjects seem to form first order beliefs relatively well.

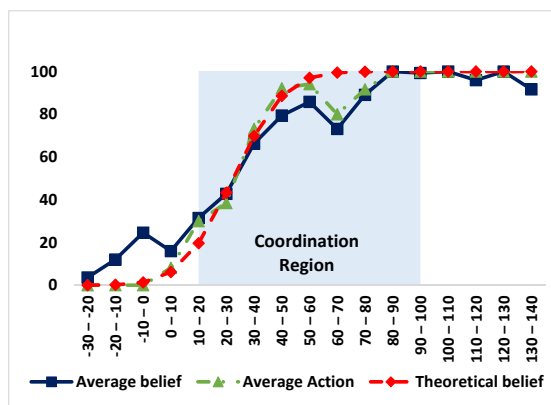
⁴⁹One might ponder that subjects can exhibit base rate neglect and that this could explain some of the observed departures, but this is not the case. As Table 15 in the Appendix shows, even if all players neglected the prior the corresponding thresholds would still be increasing in the precision of signals and converge to the risk-dominant equilibrium from below.

5.3.3 Perception of strategic uncertainty

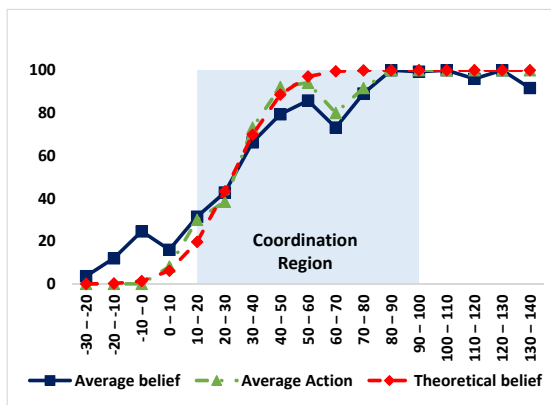
In this section, we study second order beliefs to test whether sentiments about the behavior of others can explain the results presented in Section 4.



(a) Panel A: High precision ($\sigma = 1$)



(b) Panel B: Medium precision ($\sigma = 10$)



(c) Panel C: Low precision ($\sigma = 20$)

Figure 4: Average beliefs about the action of the opponent and theoretical probabilities, by signal realization.

In Figure 4 we compare the elicited and theoretical beliefs about the probability of the opponent taking the risky action for a given signal. For illustration purposes, we divide the space of private signals into intervals of length 10 and calculate the average belief reported for the signals that fall into each interval (blue solid line).⁵⁰ The theoretical (Bayesian) beliefs are portrayed in the red

⁵⁰We do this because we observe finitely many realizations of signals.

dashed line. We also plot the average action of the subjects that observed those signals, which approximates the true probability of taking the risky action (dashed line with triangles). That is, we look at the actions taken by the subjects that observe the signals in each interval and calculate the frequency with which they, as a group, took the risky action.⁵¹ This line shows that the beliefs of subjects are consistent with their actions.

The graphs in Figure 4 support our hypothesis. Average beliefs are in general higher than equilibrium beliefs for a high signal precision. The opposite is true for low signal precision.⁵² Thus, the significantly lower thresholds associated with a high precision in Table 7 can be explained by the positive sentiments that reflect optimistic beliefs about the intention of the opponent to coordinate (left panel of Figure 4). In particular, at the estimated threshold of 19.24 we can see that subjects assign a significantly higher probability to their opponent coordinating than what the theory suggests, thus rationalizing this behavior. Likewise, the significantly higher thresholds associated with a low precision can be explained by negative sentiments that translate into pessimistic beliefs about the intention of the opponent to coordinate (right panel of Figure 4). Notice that elicited beliefs are more aligned to equilibrium beliefs for medium precision (center panel of Figure 4). This is also consistent with the estimated thresholds of Table 7, since the mean thresholds of subjects in this treatment are not statistically different from the equilibrium prediction. The distributions of elicited and equilibrium beliefs are statistically different to each other to the 1% level of significance using a Kolmogorov-Smirnov test, for all levels of precision. Finally, we note that the beliefs reported by subjects are very close to the true actions observed in the experiment suggesting that the behavior that we observed can be rationalized in the context of our extended model where beliefs are correct. Figure 10 in the appendix reports the same graphs as Figure 4 for rounds 11-50 (as opposed to rounds 26-50 as in Figure 4) to show that this sentiment-based behavior is present early in the experiment, so it is not a product of excessive learning.⁵³

Therefore, the results of these additional sessions are in line with the predictions of our extended model and provide evidence of a biased perception of strategic uncertainty. In particular, we show that sentiments switch from negative to positive as we move from an environment with high fundamental uncertainty to one with low fundamental uncertainty.

⁵¹If a subject took the safe action we assign a 0 and if he took the risky action we assign a 1. We then average these numbers across all the signals that fall into each interval and get a number between 0 and 1 that represents the true probability with which subjects took the risky action.

⁵²The high reported beliefs for negative signals are mainly due to a low number of observations and subjects who observe low signal reporting a 50-50 chance of their opponent taking either action for any signal realization in the treatment with low precision. This supports our hypothesis that a high fundamental uncertainty leads to the perception of high strategic uncertainty.

⁵³In Figure 10 we do not take into account the first 10 rounds because there is a lot of noise in the data, possibly due to subjects being acquainted with the interface and understanding the game.

5.3.4 Decreasing dispersion of thresholds

Our extended model is also consistent with an additional experimental result: a decrease in the dispersion of thresholds as information becomes more precise. In particular, note that Proposition 3 implies that as the noise in the signals decreases we should see the dispersion of thresholds decrease, since thresholds tend to the same limit as the signal noise vanishes. Tables 6 and 7 indicate that indeed we observe such a pattern in our experimental data. This provides additional validation of our extended theoretical model.

5.4 Discussion of alternative models

We have argued that the departures from the theoretical predictions can be explained by subjects having biased beliefs (or sentiments) about the behavior of their opponents. However, one may wonder whether other commonly used models of bounded rationality in games can account for the behavior of thresholds we observe in our experiment. In Section *B* of the Appendix, we discuss several popular bounded rationality models such as cursed equilibrium (Eyster and Rabin, 2005), quantal response equilibrium (see e.g., Goeree, Holt, and Pfafrey, 2016), analogy-based equilibrium (Jehiel, 2005; Koessler and Jehiel, 2008), as well as level- k and cognitive hierarchy models (see e.g., Nagel, 1995; Kneeland, 2016). We also discuss how the theoretical results are affected by assuming that agents are risk-averse, as opposed to risk neutral, and whether such an extension can help to reconcile the model with the data.

As we show in the Appendix, none of these models can explain the observed departures from the theory; at least, not without additional strong assumptions. The issue is that many of these models still predict that as information becomes precise the equilibrium converges to the risk-dominant equilibrium. The level- k model could potentially explain our results, but only if one assumes a particular relationship between signal precisions and the primitives of the level- k belief hierarchy, which is hard to justify.

6 Conclusions

In this paper we have studied how changes in the information structure affect behavior in global games. We identify two main departures from the theoretical predictions: (1) the comparative statics of thresholds with respect to signal precisions are reversed, and (2) as the signal noise decreases, subjects' behavior tends towards the efficient threshold, not the risk-dominant one. These results are starker when precisions are endogenously chosen by subjects and when beliefs are elicited. Our analysis suggests that subjects' perception of strategic uncertainty is determined by sentiments that anchor this perception to the level of fundamental uncertainty in the environment. This is

in contrast to the mechanism of the theory where strategic uncertainty increases for intermediate signals as fundamental uncertainty decreases.

To reconcile our findings with the theory, we propose a model where players can form biased beliefs about the action of their opponent, that is, they can have a biased perception of strategic uncertainty that reflects sentiments. This extended model can predict the type of thresholds that we observe in the experiment. The noise of the private signals affects sentiments by making subjects optimistic about the desire of their opponent to coordinate when information is very precise, and pessimistic when the signal noise increases. We test the mechanism of this extended model by eliciting subjects' first and second order beliefs and find evidence that subjects' beliefs are indeed biased in this way.

Our results document how a bias in belief formation can crucially affect outcomes in a game. In particular, the optimism/pessimism bias that we identify and characterize in this paper is different to the one typically studied in the behavioral literature that focuses mainly on individual decision making environments. The bias we identify is intrinsic to strategic environments, since it affects the perception of strategic uncertainty. Moreover, we show how the extent of this bias depends on the informativeness of the environment. This characterization of the bias, however, is consistent with the findings in the psychology literature if we interpret situations of high uncertainty as being "harder". In particular, the negative sentiments that arise under high fundamental uncertainty are reminiscent of the underconfidence and pessimism that arise in individual decision making when tasks are harder, and the positive sentiments under low uncertainty are similar to the overconfidence or optimism that arise in individual decision making for easier tasks (see e.g., Moore and Cain (2007)).

The departures from the theory that we identify have important welfare considerations. Our extended model and the evidence that supports it propose a novel mechanism to understand behavior in environments characterized by strategic complementarities and incomplete information.

Our results can also shed some light on recent stylized facts in macroeconomics. For example, our mechanism explains how recessions can be associated with an increase in uncertainty by noticing that an increase in uncertainty leads to negative sentiments about the likelihood of a profitable risky investment (via pessimism about others investing). This leads to lower levels of investment, which amplify a recession. This observation reconciles the views of Bloom (2009) and Angeletos and La'O (2013) and Benhabib et al. (2015). In terms of policy making, our results suggest that greater transparency in financial regulation might be beneficial in environments with strategic complementarities, like bank runs. In this case, an increase in transparency can lead to positive sentiments about the likelihood of others rolling over their loans, which leads to less early withdrawals and greater financial stability.

Appendix

A Additional experimental results

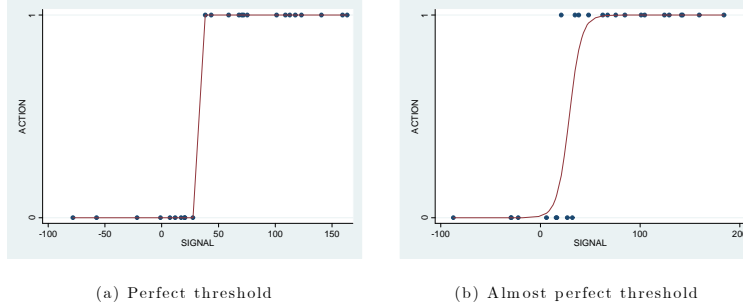


Figure 5: Examples of perfect and almost perfect thresholds.

	Prec 1	Prec 2	Prec 3	Prec 4	Prec 5	Prec 6
Prec 1	0.95	0.03	0	0	0	0.02
Prec 2	0.08	0.74	0.08	0.05	0	0.05
Prec 3	0	0.02	0.87	0.09	0	0.02
Prec 4	0	0	0.04	0.92	0.02	0.02
Prec 5	0.01	0.01	0.10	0.15	0.58	0.15
Prec 6	0.01	0	0.02	0.04	0.03	0.90

Figure 6: Transition matrix of precision choices in the last 25 rounds, endogenous information.

Precision	High	Medium	Low
Realized payoffs	34.54	33.3	29.71
	(8.82)	(13.48)	(9.65)
Expected equilibrium payoffs	23.62***	29.1***	31.44***
	(8.77)	(10.42)	(5.97)
Expected constrained efficient payoffs	29.26***	31.58***	31.68***
	(7.11)	(10.83)	(5.96)
First-best complete information payoffs	36.71***	39.46***	39.8***
	(6.4)	(10.89)	(4.56)

Different from realized payoffs at the ***1%; **5%; *10% level

Table 9: Average payoffs and efficiency benchmarks, exogenous information.

Variable	Action $\{A,B\} = \{0,1\}$
Precision 1*signal	0.178*** (0.03)
Precision 2*signal	0.191*** (0.05)
Precision 3*signal	0.096*** (0.01)
Precision 4*signal	0.088*** (0.01)
Precision 5*signal	0.062*** (0.02)
Precision 6*signal	0.057*** (0.01)
Constant	-3.559*** (0.46)
N	1000

Clustered (by subject) standard errors in parentheses; * significant at 10%; ** significant at 5%; *** significant at 1%

Table 10: Random effects logit: risky action as a function of precision, endogenous information and direct action choice.

Variable	Reported threshold
Precision 2	6.25 (4.00)
Precision 3	10.65*** (3.48)
Precision 4	10.78*** (3.39)
Precision 5	12.87*** (3.66)
Precision 6	12.54*** (3.23)
Constant	20.66*** (5.14)
N	1100

Clustered (by subject) standard errors in parentheses; * significant at 10%; ** significant at 5%; *** significant at 1%

Table 11: Random effects OLS: reported threshold as a function of precision, endogenous information and strategy method

		Choice of pair member 2					
		Prec 1	Prec 2	Prec 3	Prec 4	Prec 5	Prec 6
Choice of pair member 1	Prec 1	7.24%	2.19%	5.71%	4.29%	1.24%	1.52%
	Prec 2		0.57%	0.29%	3.52%	0.10%	0.19%
	Prec 3			5.43%	14.00%	0.57%	5.33%
	Prec 4				20.57%	1.71%	9.05%
	Prec 5					0.10%	4.00%
	Prec 6						12.38%

Table 12: Combination of precision choices, endogenous information.

Precision	High	Medium	Low
Realized payoffs	32.94 (7.25)	26.16 (8.61)	18.11 (12.97)
Expected equilibrium payoffs	31.06*** (6.65)	30.83*** (6.66)	33.63*** (5.75)
Expected constrained efficient payoffs	31.21*** (6.83)	30.97*** (6.84)	33.84*** (5.90)
First-best complete information payoffs	35.25*** (6.61)	35.02*** (6.62)	37.80*** (5.71)

Different from realized payoffs at the ***1%; **5%; *10% level

Table 13: Average payoffs and efficiency benchmarks, endogenous information.

Evolution of thresholds over time. Figure 7 shows the average difference in absolute value between the threshold that a subject reports in one period with respect to his threshold in the previous period, for each precision level when information is exogenous and for subjects in pairs that coordinate on high, medium, and low precisions when information is endogenous. The vertical bar at period t illustrates how much, on average, a subject changes the value of his own threshold in period t with respect to the threshold he reported in period $t - 1$. Figure 8 illustrates convergence within a pair by plotting the average difference in absolute value between the reported thresholds of both members of each pair in each period, for each precision level with exogenous and endogenous information structures. The vertical bar at period t illustrates how much, on average, subjects coordinate their actions with their pair member during that period.

Additional analysis: weights given to private signals under belief elicitation. We assume that beliefs are linear in the prior mean and in their signal (as Bayes' rule prescribes) but we allow for the possibility that players use non-Bayesian weights. In particular, we postulate that subject i 's stated belief θ_i^B satisfies $\theta_i^B = w_i x_i + (1 - w_i) \mu_\theta$. According to Bayes' rule, the correct weight of the private signal is $w_x^{Bayes} = \frac{\sigma^{-2}}{\sigma^{-2} + \sigma_\theta^{-2}}$. To estimate w_i for each subject i , using the above formula, we back out these weights for the last 25 rounds. We then take the average of the implied

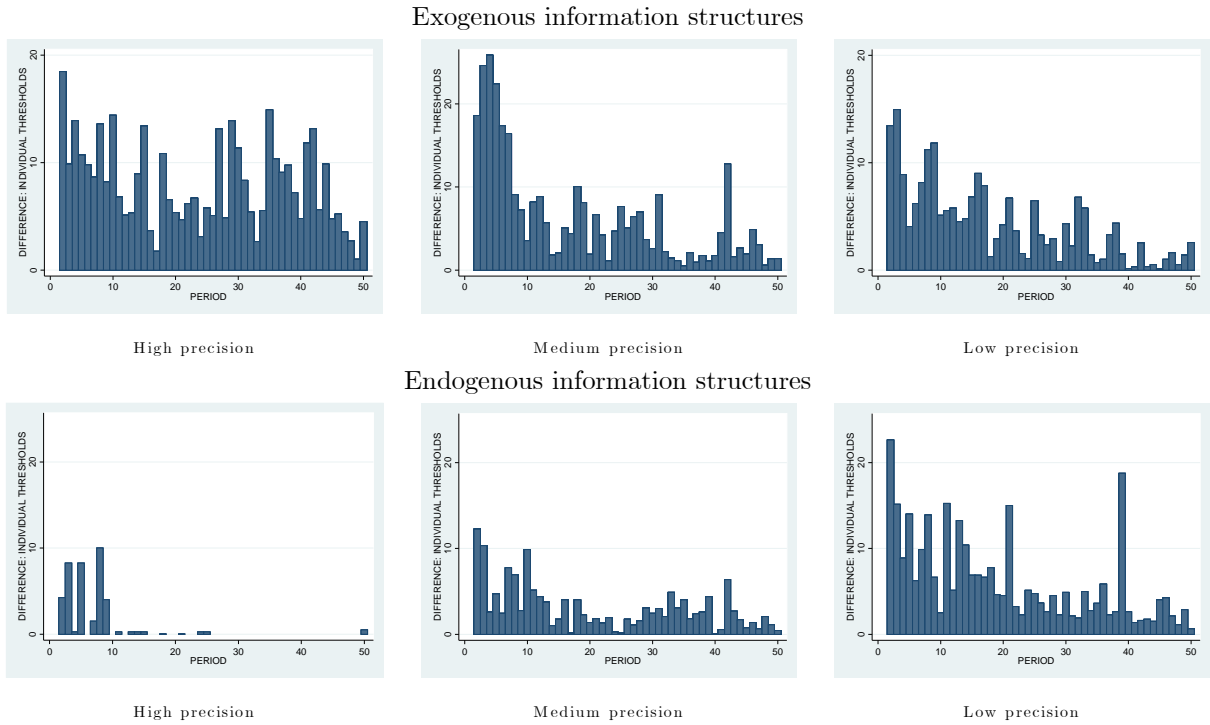


Figure 7: Convergence of individual thresholds, strategy method treatments.



Figure 8: Convergence of thresholds within pairs, strategy method treatments.

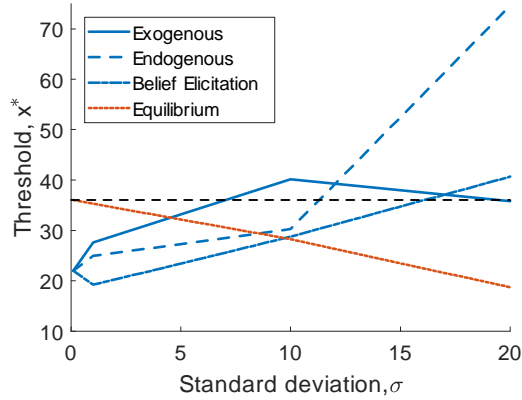


Figure 9: Theoretical predictions and estimated thresholds for exogenous and endogenous information structures

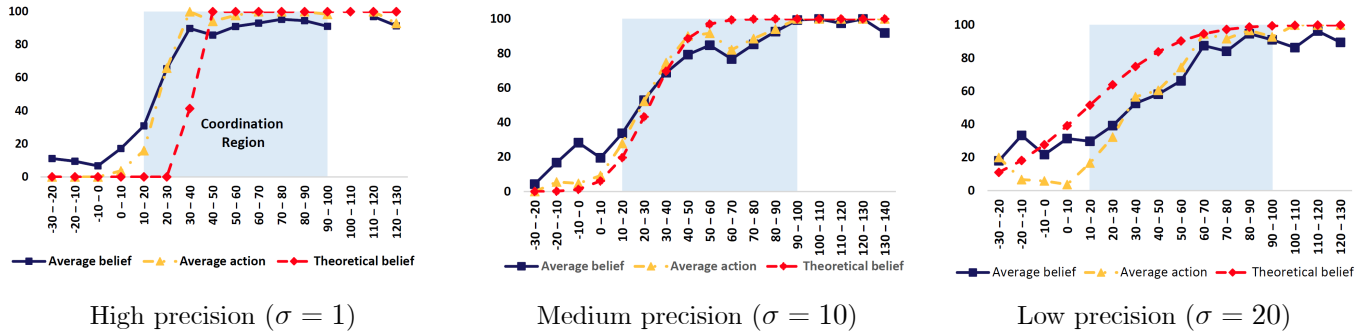


Figure 10: Average beliefs about the action of the opponent and theoretical probabilities for rounds 11-50, by signal realization.

weights to obtain subject i 's average weight w_i .⁵⁴

⁵⁴To compute average weight w_i we discard the rounds in which the estimated weights w_{it} are lower than -2 or larger than 3 . The interval $[-2, 3]$ contains more than 95% of all estimated w_{it} 's (and virtually all for the case of high precision). This eliminates the weights that seem to be a result of mistakes (e.g., reporting a belief of -586 when signal was -4.02). Note that a weight above 1 implies that a subject updated his belief towards the signal, but his reported belief lies further from the prior mean than the signal. On the other hand, a weight lower than 0 implies that a subject updated his belief away from the private signal. Typically, we observed the weights to fall outside the interval $[0, 1]$ when signals were close to the mean of the prior (for example, the signal is 49 and the stated belief is 52, resulting in a weight of -2). While these beliefs are hard to justify from a Bayesian point of view, they are not completely unreasonable.

Precision	High	Medium	Low
w_i^{Bayes}	0.99	0.96	0.86
Mean (w_i)	0.98	0.97	0.93
Median (w_i)	1.00	0.94	0.98
St. dev. (w_i)	0.04	0.12	0.17

Table 14: Estimated weights given to private signals, by precision.

Precision	High	Medium	Low
Thresholds with base rate neglect	35.43	30.36	24.65
Thresholds without base rate neglect	35.31	28.31	18.73

Table 15: Thresholds predicted by the theory when players neglect the prior and when they do not, by precision

B Alternative models and our data

In this Section we discuss some popular bounded rationality models and show that these models have a hard time explaining our experimental findings.

Cursed equilibrium One equilibrium concept that incorporates bounded rationality similarly to us is the concept of Cursed Equilibrium (Eyster and Rabin, 2005). This model also relaxes the assumption that players form correct beliefs about their opponents, and does so in a very particular way. Let λ measure the degree of “cursedness” of a player. With probability λ , a player believes that his opponent randomizes his action choice uniformly and with the remaining probability he forms Bayesian beliefs. Therefore, $\lambda = 0$ corresponds to full rationality and $\lambda = 1$ to full cursedness.

This specific departure from Bayesian beliefs cannot organize our findings. In our game, $\lambda = 0$ corresponds to the standard model where equilibrium thresholds converge to the risk-dominant threshold as $\sigma \rightarrow 0$. On the other hand, with $\lambda = 1$ the resulting equilibrium thresholds would be even closer to the risk-dominant threshold of $2T$ since in this case players expect their opponents to randomize their actions uniformly.⁵⁵ As a consequence, regardless of what value λ takes and whether it varies with the information structure, cursed equilibrium predicts that the thresholds will converge to the risk dominant equilibrium as the noise in the private signals vanishes, which is inconsistent with our experimental findings.

⁵⁵Best responding to full cursedness is equivalent to the best response in a risk dominant equilibrium, which corresponds to the equilibrium in global games in the limit as the signal noise vanishes (see Carlsson and Van Damme, 1993, and Morris and Shin, 2003).

Quantal response equilibrium Another alternative equilibrium concept that is widely used to organize experimental data is quantal response equilibrium (QRE). Unlike our extended model that allows for “mistakes” in belief formation, QRE assumes that players form correct beliefs, but that they sometimes make mistakes when choosing an action.⁵⁶ According to QRE, deviations from the Nash Equilibrium actions are less likely to occur when the cost of taking the wrong action increases.

Following Goeree, Holt, and Palfrey (2016), let λ measure the responsiveness of a player to expected payoffs. When $\lambda = 0$ a player does not respond to expected payoffs and randomizes between actions, and when $\lambda \rightarrow \infty$ the player plays a best-response consistent with Nash Equilibrium. Define $\omega_{iA} = F(\lambda(U_{iA} - U_{iB}))$ as the probability with which player i chooses action A , where F corresponds to a strictly increasing continuously differentiable CDF, U_{iA} is player i 's expected payoff of taking action A and U_{iB} is player i 's expected payoff of taking action B .

While QRE can explain some characteristics of the thresholds we observe (such as the observed frequency of perfect and almost perfect thresholds across precision levels) it is unlikely to explain why the observed thresholds tend to decrease and converge to the efficient equilibrium as information becomes more precise. This is because choosing the “wrong” action is more costly, in terms of differences in expected utility, when signals are precise since in this case players can predict well whether the risky action will be successful or not.⁵⁷ As such, QRE would suggest that players will use similar thresholds to those predicted by standard global games when information is very precise. Thus, QRE by itself cannot explain why, as the signal precision increases, the estimated thresholds decrease and seem to converge to the efficient threshold.

Level- k reasoning Another possible way to explain our findings is to think about how different signal precisions can affect the level of reasoning of subjects according to models like level- k and cognitive hierarchy. These models assume that players have limited depths of reasoning, referred to as “levels,” with a player who can perform k levels of reasoning called a level- k type. Level-0 are assumed to be non-strategic types who randomize their actions according to a pre-specified distribution, and all players of level- k with $k > 1$ best respond to their exogenous belief distribution over lower types (see Strzalecki (2014) for a general formulation of this model).

⁵⁶In particular, our extended model relaxes rationality in belief formation and assumes that players best-respond to beliefs, while QRE assumes that players form correct beliefs and relaxes the assumption of best-responses.

⁵⁷To see how QRE can explain the observed frequency of perfect and almost perfect thresholds across precision levels note that as signals get closer to the equilibrium threshold players grow uncertain about whether the risky action will be successful or not. For these signals, the expected cost of making a mistake is larger as signals become more precise because subjects' beliefs about the state become more accurate and correlated within a pair. In contrast, for a low precision the expected cost of making a mistake for signals close to the threshold is not as large since there is more uncertainty about the fundamentals and about the beliefs of the opponent, since signals are not very correlated. This implies that we should see a clear switching point for subjects under high precision (a perfect threshold in Figure 5) and we should see subjects alternating between actions for the signals closer to the threshold under a low precision (an almost-perfect threshold in Figure 5), which is indeed the case.

In the context of global games, Kneeland (2016) shows that the results from HNO04 are more in line with a model of limited depth of reasoning than with equilibrium play. In particular, level- k explains the use of threshold strategies under complete information (as is found in our data). However, in order for this model to organize our finding about reversed comparative statics of thresholds with respect to precisions and to generate the observed distribution of second-order beliefs, one would need to assume that both the behavior of level-0 players and the distribution of types vary with the signal precision.⁵⁸ Thus, while the level- k model could potentially explain our results, it requires strong assumptions about the relationship between signal precision and the primitives of the level- k belief hierarchy. In contrast, our extended model has an intuitive interpretation that does not require strong assumptions on primitives to get predictions consistent with the data.

Analogy-based expectations Jehiel (2005) and Koessler and Jehiel (2008) develop the concept of “analogy-based equilibria” in games. They argue that it is implausible to assume that agents understand their opponents’ strategy state-by-state. Instead, they suggest that players, when forming their beliefs about other players’ strategies, bundle together nearby state into bins (referred to as analogy partitions) and best respond to the opponents’ average action in each bin. The analogy-based equilibrium has intuitive appeal and, as shown by Koessler and Jehiel (2008), it can be applied to global games. However, our findings cannot be organized by this concept either.

To see this, and following Koessler and Jehiel (2008), assume that subjects use the coarsest analogy partition. This means that they best-respond to the average behavior of their opponent. Thus, based on his signal, player i ’s expects that player j will choose action A with some probability P_j , where P_j is the actual ex-ante probability of agent j choosing A given player j ’s strategy. Thus, player i will choose A if and only if

$$E[\theta|x_i] P_j - T \geq 0$$

From the above equation, it is immediate that players will follow threshold strategies. Therefore,

⁵⁸ Alaoui and Penta (2015) interpret the level- k model as an optimization problem where players perform a cost-benefit analysis to determine the level of reasoning they use in the game. In the context of our experiment, one could imagine that different precisions can affect the distribution of levels played because the signal precision affects the expected payoff, and thus the benefit of taking the risky action. Even if this can imply a different proportion of level-0 players for different precisions, it does not affect the way in which a level-0 player behaves. If level-0 behavior depended on the informativeness of the environment, then level-0 players would be somewhat strategic, thus violating the main postulate in the literature about level-0 players being non-strategic.

the equilibrium is given by the following two equations:

$$\begin{aligned} E[\theta|x_i^*] \Pr(x_j \geq x_j^*|x_i^*) - T &= 0 \\ E[\theta|x_j^*] \Pr(x_i \geq x_i^*|x_j^*) - T &= 0 \end{aligned}$$

Any pair of thresholds that $\{x_i^*, x_j^*\}$ that satisfies the above equations is an equilibrium. Note, however, that as $\sigma \rightarrow 0$ we have that

$$\Pr(x_j \geq x_j^*|x_i^*) \rightarrow \frac{1}{2} \quad \text{and} \quad E[\theta|x_i^*] \rightarrow x_i^*$$

Therefore, in the limit, the equilibrium still converges to the risk-dominant equilibrium. Thus, the analogy-based equilibrium cannot explain the fact that threshold converge towards the efficient equilibrium as noise vanishes.

Risk-averse agents In the model we assumed that agents are risk-neutral. Is it possible that adding risk aversion could help to explain our experimental finding? Assume that players have utility $u(w) = w^\rho$, where w is the amount they earn in a given round and $\rho \in (0, 1)$. Let E denote the endowment of a player (recall that in the experiment agents start with a positive endowment). If action A is successful then the payoff is $\theta + E - T$. If action A is unsuccessful, then the payoff is $E - T$.⁵⁹ In that case, the equilibrium thresholds solve

$$\begin{aligned} &\int_{\underline{\theta}}^{\bar{\theta}} (\theta + E - T)^\rho \Pr(x_j \geq x^*|\theta) f(\theta|x^*) d\theta + \int_{\underline{\theta}}^{\infty} (\theta + E - T)^\rho f(\theta|x^*) d\theta \\ &+ \int_{\underline{\theta}}^{\bar{\theta}} (E - T)^\rho \Pr(x_j < x^*|\theta) f(\theta|x^*) d\theta + \int_{-\infty}^{\underline{\theta}} (E - T)^\rho f(\theta|x^*) d\theta \\ &= E \end{aligned}$$

The first two terms on the LHS of the above equation correspond to the case where action A is successful, in which case player i earns $\theta + E - T$ if he chooses action A . The last two terms on the LHS correspond to the case when action A is unsuccessful, in which case player i receives payoff $E - T$ if he chooses action A . The payoff from choosing action B is E .

Let $\sigma \rightarrow 0$. In that case, the above equation converges to

$$\frac{1}{2} (x_i^* + E - T)^\rho + \frac{1}{2} (E - T)^\rho = E$$

⁵⁹When players are risk-neutral the endowment does not affect their choices.

Therefore, in the limit

$$\frac{\partial x_i^*}{\partial \rho} = -\frac{\frac{1}{2} [\ln(x_i^* + E - T)] (x_i^* + E - T)^\rho + \frac{1}{2} [\ln(E - T)] (E - T)^\rho}{\rho^{\frac{1}{2}} (x_i^* + E - T)^{\rho-1}} < 0$$

Thus, the more risk-averse are the players (i.e., the lower ρ) the higher is the threshold. Thus, risk aversion would push thresholds above the risk-dominant threshold in the limit as information becomes precise. It can be shown that in order to match our findings, subjects would have to be risk-loving when information is precise and risk averse when information is imprecise, which is hard to justify.

C Baseline model with heterogenous precisions - FOR ONLINE PUBLICATION

In this section we provide the proofs of the claims stated in Section 2. We first focus on the global game with heterogenous precisions and then extend it to the two stage game with costly information acquisition.

C.1 Model with exogenous information

C.1.1 Relation to monotone supermodular games

We first prove that the coordination game with heterogenous precisions belongs to the class of monotone supermodular games as defined by Vives and Van Zandt (2007). Following their notation, define $N = \{1, 2\}$ as the set of players indexed by i . Let the type space of player i be a measurable space $(\Omega_i, \mathcal{F}_i)$. Denote by $(\Omega_0, \mathcal{F}_0)$ a state space that is capturing the residual uncertainty.⁶⁰ We let \mathcal{F}_{-i} be the product σ -algebra $\otimes_{k \neq i} \mathcal{F}_k$. Let player i 's interim beliefs be given by a function $p_i : \Omega_i \rightarrow M_{-i}$, where M_{-i} is the set of probability measures on $(\Omega_{-i}, \mathcal{F}_{-i})$. Finally, let $A_i = \{0, 1\}$ be the action set of player i , A be the set of action profiles and $u_i : A \times \Omega \rightarrow \mathbb{R}$ be the payoff function.

Definition 4 *A game belongs to the class of monotone supermodular games if*

1. *The utility function $u_i(a_i, a_{-i}, \omega)$ is supermodular in own actions, a_i , and has increasing differences in (a_i, a_{-i}) and in (a_i, ω) .*
2. *The belief map $p_i : \Omega_i \rightarrow M_{-i}$ is increasing with respect to a partial order on M_{-i} of first-order stochastic dominance.*

⁶⁰In a global games setting we usually interpret $(\Omega_i, \mathcal{F}_i)$ to be the space of possible signals that agent i receives, while $(\Omega_0, \mathcal{F}_0)$ corresponds to the measurable space of the underlying parameters of the game.

In our case, the type space is defined as follows: $\Omega_0 = \mathbb{R}$, $\Omega_i = \mathbb{R}$ for $i = 1, 2$, where $\omega_0 = \theta$, $\omega_i = \hat{\theta}_i$, $\omega_j = \hat{\theta}_j$ and $\mathcal{F}_i = B(\mathbb{R})$, a Borel σ -algebra on \mathbb{R} , $i = 0, 1, 2$. The set of probability measures M_{-i} is simply the set of joint normal probability distributions over $(\Omega_{-i}, \mathcal{F}_{-i})$ conditional on the realization of ω_i . The belief mapping $p_i : \Omega_i \rightarrow M_{-i}$ maps $\hat{\theta}_i$ into the posterior distribution of $(\theta, \hat{\theta}_j)$ using Bayes' rule. Finally, the underlying utility function for player i is given by

$$u(a_i, a_j, \theta) = 1_{\{a_i=1\}} \left[\theta \left[1_{\{\theta \in [\underline{\theta}, \bar{\theta}]\}} 1_{\{a_j=1\}} + 1_{\{\theta > \bar{\theta}\}} \right] - T \right]$$

and the expected utility of player i is:

$$v_i(a_i, a_j, \theta) = 1_{\{a_i=1\}} \left[\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \theta \left[1_{\{\theta \in [\underline{\theta}, \bar{\theta}]\}} 1_{\{s_j(\hat{\theta}_j)=1\}} + 1_{\{\theta > \bar{\theta}\}} \right] \frac{1}{\hat{\sigma}_i} \frac{1}{\sigma_j} \phi \left(\frac{\theta - \hat{\theta}_i}{\hat{\sigma}_i} \right) \phi \left(\frac{\sigma_j (\hat{\theta}_j - \tilde{\theta}_j)}{\hat{\sigma}_j^2} \right) d\hat{\theta}_j d\theta \right] - T$$

where $s_j : \Omega_j \rightarrow A_j$ is a measurable strategy of player j .

We proceed now by extending the following result from Van Zandt and Vives (2007) for unbounded utility functions.

Proposition 5 (van Zandt and Vives) *Assume that a game Γ belongs to the class of monotone supermodular games. Furthermore, assume that the following hold:*

1. Each Ω_k is endowed with a partial order,
2. A_i is a complete lattice,
3. $\forall a_i \in A_i$, $u_i(a_i, \cdot) : \Omega \rightarrow \mathbb{R}$ is measurable,
4. u_i is bounded.
5. u_i is continuous in a_i .⁶¹

Then, there exist a least and a greatest Bayesian Nash Equilibrium of the game Γ and each one of them is in monotone strategies.

The fact that global games in general belong to the class of monotone supermodular games was noted first by Van Zandt and Vives (2007).

It is straightforward to see that the game with heterogenous precisions also belongs to this class of games. However, the utility function in our model does not satisfy the restrictions imposed in Proposition 5 as it is unbounded.⁶² We now show that the result of Vives and Van Zandt (2007) can be extended to integrable, potentially unbounded utility functions.

⁶¹When A_i is finite this condition is vacuous.

⁶²However, note that u is bounded from below by $-(|\theta| + T)$ which is integrable with respect to the measure μ_F implied by a player's posterior belief.

Proposition 6 *Assume that the game Γ belongs to the class of monotone supermodular games. Furthermore, assume that conditions (1) – (3) of Proposition 1 are satisfied, and that u satisfies the following assumption:*

(1C) *There exists a measurable function h that is integrable with respect to $p(\omega_{-i}|\omega_i)$ for all ω_i , all ω_{-i} , and $|u| < h$.*

Then there exists a least and a greatest Bayesian Nash Equilibrium of the game Γ and each one of them is in monotone strategies.

Proof. We prove the above result in two steps. First, assuming that the greatest best reply mapping $\bar{\beta}_i$ is well-defined, increasing, and monotone, we show that the greatest Bayesian Nash Equilibrium (BNE) exists. Then, we show that under the above conditions $\bar{\beta}_i$ is indeed well-defined, increasing, and monotone.

Step 1: Suppose that $\bar{\beta}_i$ is well-defined, increasing and monotone and u satisfies assumption (1C). Then we can repeat the argument of Van Zandt and Vives (2007) to show that there is a greatest and least BNE in monotone strategies. We can relax the boundedness assumption, since under assumption (1C) we can interchange the order of limit and integration invoking the Lebesgue Dominated Convergence Theorem. Since this is the only step in that proof that requires boundedness of the utility function, we are done.

Step 2: Here we need to establish that $\bar{\beta}_i$ is well-defined and increasing. Then, the monotonicity of $\bar{\beta}_i$ will follow from Proposition 11 in Van Zandt and Vives (2007). The only difficult part of this step is to show that $\bar{\beta}_i$ is well-defined, and more precisely that it is a measurable function of ω_i . For this purpose we extend the proof of Lemma 9 in Ely and Peski (2006) to cover more general measurable functions. The rest of argument follows from Van Zandt (2010).

Fix $a_i \in A_i$ and define $U_i(\omega_i, \omega_j) := u_i(a_i, s_j(\omega_j), \omega_i, \omega_{-i})$. We need to show that a function $\pi_i : A_i \times \Omega_i \rightarrow \bar{\mathbb{R}}$ defined by

$$\pi_i(a_i, \omega_i) = \int_{\Omega_{-i}} U_i(\omega_i, \omega_{-i}) dp(\omega_{-i}|\omega_i)$$

is measurable in t_i . To prove this we use a result by Ely and Peski (2006):

Lemma (Ely and Peski) *Let A and B be measurable sets and $g : A \times B \rightarrow [0, 1]$ be a jointly measurable map. If $m : A \rightarrow \Delta B$ (where ΔB denotes the set of probability measures defined on B) is measurable, then the map $L^g : A \rightarrow \mathbb{R}$ defined as $L^g(a) = \int g(a, \cdot) dm(a)$ is measurable.*

Note however, that the proof of their lemma is essentially unchanged if we allow $g : A \times B \rightarrow \bar{\mathbb{R}}$, as long as g is integrable and bounded from below by an integrable function h . In this case, there exists a sequence of simple functions g_n such that $g_n \rightarrow g$, so by the extended Monotone Convergence Theorem (Ash, 2000) we have $\int g_n dv \rightarrow \int g dv$ for a measure v defined on $A \times B$.

Hence we conclude that $\pi_i : A_i \times \Omega_i \rightarrow \overline{\mathbb{R}}$ is a measurable function of ω_i . The rest of the proof follows directly from Van Zandt (2010) Section 7.5. Monotonicity of $\overline{\beta}_i$ follows from Proposition 11 in Van Zandt and Vives (2007). ■

C.1.2 Equilibrium uniqueness

In order to facilitate notation when solving the model, we rewrite the condition for threshold strategies in terms of the posteriors that players hold about the fundamental θ , as in Hellwig (2002). Note that this is straightforward since the posterior about θ held by player i , $\widehat{\theta}_i$, is a linear, strictly increasing function of the signal he observes, x_i . Therefore, player i will take the risky action whenever his posterior belief about θ , given his signal realization of x_i , is higher than the posterior of θ that corresponds to player i 's optimal threshold:

$$a(x_i; \sigma) = \begin{cases} 1 & \text{iff } \widehat{\theta}_i \geq \widehat{\theta}_i^*(\sigma) \\ 0 & \text{iff } \widehat{\theta}_i < \widehat{\theta}_i^*(\sigma) \end{cases}$$

where $\widehat{\theta}_i^* = \frac{\mu_\theta \sigma_i^2 + x_i^* \sigma_\theta^2}{\sigma_i^2 + \sigma_\theta^2}$. In order to write the condition of player j in terms of his posterior belief, notice that

$$\frac{x_j^* - \theta}{\sigma_j} = \frac{\sigma_j (\widehat{\theta}_j^* - \widetilde{\theta}_j)}{\widehat{\sigma}_j^2}$$

$$\text{where } \widetilde{\theta}_j = \frac{\sigma_\theta^2 \theta + \sigma_j^2 \mu_\theta}{\sigma_j^2 + \sigma_\theta^2}.$$

The expected payoff of taking the risky action for player $i = 1, 2$, conditional on observing signal x_i and given that the other player follows a threshold strategy with switching point $\widehat{\theta}_j^*$ is:

$$v_i(x_i, x_j^*; \sigma) = \frac{1}{\widehat{\sigma}_i} \int_{\underline{\theta}}^{\bar{\theta}} \theta \phi \left(\frac{\theta - \widehat{\theta}_i}{\widehat{\sigma}_i} \right) \left(1 - \Phi \left(\frac{\sigma_j (\widehat{\theta}_j^* - \widetilde{\theta}_j)}{\widehat{\sigma}_j^2} \right) \right) d\theta + \frac{1}{\widehat{\sigma}_i} \int_{\bar{\theta}}^{\infty} \theta \phi \left(\frac{\theta - \widehat{\theta}_i}{\widehat{\sigma}_i} \right) d\theta - T \quad (6)$$

This notation is used for all proofs in this appendix.

Lemma 1 *The payoff for player i of taking the risky action, $v_i(x_i, x_j^*; \sigma)$, is increasing in his own signal x_i , and decreasing in the other player's threshold x_j^* , for $i, j = 1, 2$, $i \neq j$.*

Proof. (1) Note that $\widehat{\theta}_i$ is an increasing function of x_i , i.e. $\frac{\partial \widehat{\theta}_i}{\partial x_i} > 0$. Thus, it is enough to show that the payoff of taking the risky action is increasing with respect to the posterior mean of θ , $\widehat{\theta}_i$.

Taking a partial derivative of (6) wrt $\widehat{\theta}_i$ yields:

$$-\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\widehat{\sigma}_i^2} \phi' \left(\frac{\theta - \widehat{\theta}_i}{\widehat{\sigma}_i} \right) \left(1 - \Phi \left(\frac{\sigma_j (\widehat{\theta}_j^* - \widetilde{\theta}_j)}{\widehat{\sigma}_j^2} \right) \right) d\theta - \int_{\bar{\theta}}^{\infty} \theta \frac{1}{\widehat{\sigma}_i^2} \phi' \left(\frac{\theta - \widehat{\theta}_i}{\widehat{\sigma}_i} \right) d\theta$$

Applying integration by parts to the second term of the above expression and simplifying we see that this term is equal to

$$\bar{\theta} \frac{1}{\hat{\sigma}_i} \phi \left(\frac{\bar{\theta} - \hat{\theta}_i}{\hat{\sigma}_i} \right) + \left(1 - \Phi \left(\frac{\bar{\theta} - \hat{\theta}_i}{\hat{\sigma}_i} \right) \right) > 0 \quad (7)$$

Similarly, applying integration by parts to the first term of the above derivative we obtain

$$\begin{aligned} & - \left[\theta \frac{1}{\hat{\sigma}_i} \phi \left(\frac{\theta - \hat{\theta}_i}{\hat{\sigma}_i} \right) \left(1 - \Phi \left(\frac{\sigma_j (\hat{\theta}_j^* - \tilde{\theta}_j)}{\hat{\sigma}_j^2} \right) \right) \right]_{\underline{\theta}}^{\bar{\theta}} \\ & + \int_{\underline{\theta}}^{\bar{\theta}} \frac{1}{\hat{\sigma}_i} \phi \left(\frac{\theta - \hat{\theta}_i}{\hat{\sigma}_i} \right) \left(1 - \Phi \left(\frac{\sigma_j (\hat{\theta}_j^* - \tilde{\theta}_j)}{\hat{\sigma}_j^2} \right) \right) d\theta + \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\hat{\sigma}_i \sigma_j} \phi \left(\frac{\theta - \hat{\theta}_i}{\hat{\sigma}_i} \right) \phi \left(\frac{\sigma_j (\hat{\theta}_j^* - \tilde{\theta}_j)}{\hat{\sigma}_j^2} \right) d\theta \quad (8) \end{aligned}$$

Putting (7) and (8) together we get:

$$\begin{aligned} & \bar{\theta} \frac{1}{\hat{\sigma}_i} \phi \left(\frac{\bar{\theta} - \hat{\theta}_i}{\hat{\sigma}_i} \right) + \left(1 - \Phi \left(\frac{\bar{\theta} - \hat{\theta}_i}{\hat{\sigma}_i} \right) \right) - \bar{\theta} \frac{1}{\hat{\sigma}_i} \phi \left(\frac{\bar{\theta} - \hat{\theta}_i}{\hat{\sigma}_i} \right) \left(1 - \Phi \left(\frac{\sigma_j (\hat{\theta}_j^* - \tilde{\theta}_j)}{\hat{\sigma}_j^2} \right) \right) \\ & + \theta \frac{1}{\hat{\sigma}_i} \phi \left(\frac{\theta - \hat{\theta}_i}{\hat{\sigma}_i} \right) \left(1 - \Phi \left(\frac{\sigma_j (\hat{\theta}_j^* - \tilde{\theta}_j)}{\hat{\sigma}_j^2} \right) \right) + \int_{\underline{\theta}}^{\bar{\theta}} \frac{1}{\hat{\sigma}_i} \phi \left(\frac{\theta - \hat{\theta}_i}{\hat{\sigma}_i} \right) \left(1 - \Phi \left(\frac{\sigma_j (\hat{\theta}_j^* - \tilde{\theta}_j)}{\hat{\sigma}_j^2} \right) \right) d\theta \\ & + \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\hat{\sigma}_i \sigma_j} \phi \left(\frac{\theta - \hat{\theta}_i}{\hat{\sigma}_i} \right) \phi \left(\frac{\sigma_j (\hat{\theta}_j^* - \tilde{\theta}_j)}{\hat{\sigma}_j^2} \right) d\theta > 0 \quad (9) \end{aligned}$$

since $\bar{\theta} \frac{1}{\hat{\sigma}_i} \phi \left(\frac{\bar{\theta} - \hat{\theta}_i}{\hat{\sigma}_i} \right) \geq \bar{\theta} \frac{1}{\hat{\sigma}_i} \phi \left(\frac{\bar{\theta} - \hat{\theta}_i}{\hat{\sigma}_i} \right) \left(1 - \Phi \left(\frac{\sigma_j (\hat{\theta}_j^* - \tilde{\theta}_j)}{\hat{\sigma}_j^2} \right) \right)$. Since $\hat{\theta}_i$ is an increasing linear function of x_i , this proves the first claim.

(2) Similarly, we note that the posterior $\hat{\theta}_j^*$ is strictly increasing in x_j^* . Hence, it is enough to show that the derivative of the payoff of taking the risky action for player i with respect to $\hat{\theta}_j^*$ is negative. But the partial derivative of (6) wrt $\hat{\theta}_j^*$ is given by

$$- \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{\sigma_j}{\hat{\sigma}_i \hat{\sigma}_j^2} \phi \left(\frac{\theta - \hat{\theta}_i}{\hat{\sigma}_i} \right) \phi \left(\frac{\sigma_j (\hat{\theta}_j^* - \tilde{\theta}_j)}{\hat{\sigma}_j^2} \right) d\theta < 0$$

This establishes that $v_i(x_i, x_j^*; \sigma)$ is decreasing in x_j^* . ■

Theorem 1 *There exists a unique, dominance solvable equilibrium of the coordination game in which both players use threshold strategies characterized by (x_1^*, x_2^*) if either:*

1. $\frac{\sigma_i}{\sigma_\theta} < K_i(\underline{\theta}, \bar{\theta}, \mu_\theta)$, $i = 1, 2$ holds, for any pair of (σ_1, σ_2) , or
2. $\sigma_\theta > \bar{\sigma}_\theta$, where $\bar{\sigma}_\theta$ is determined by the parameters of the model.

Proof. As proven above, the coordination game belongs to the class of monotone supermodular games and therefore we know that there are a least and a greatest Bayesian Nash Equilibria in monotone strategies. To prove the theorem, we only need to show that these equilibria are the same, i.e. that there is a unique equilibrium in threshold strategies. For ease of exposition, we perform the analysis in terms of thresholds over posterior beliefs, $(\hat{\theta}_1^*, \hat{\theta}_2^*)$. Uniqueness of these thresholds imply uniqueness of thresholds over signals (x_1^*, x_2^*) .

Let, $s_i(\hat{\theta}_i^*)$ be a threshold strategy of player i with switching point $\hat{\theta}_i^*$ such that $s_i(\hat{\theta}_i^*) = 1$ (risky action) if $\hat{\theta}_i \geq \hat{\theta}_i^*$ and $s_i(\hat{\theta}_i^*) = 0$ (safe action) if $\hat{\theta}_i < \hat{\theta}_i^*$, where $\hat{\theta}_i$ is the posterior belief that player i holds about θ after observing signal x_i , for $i = 1, 2$. Then the equilibrium conditions are given by the following equations:

$$v_1(\hat{\theta}_1^*, \hat{\theta}_2^*; \boldsymbol{\sigma}) \equiv \frac{1}{\hat{\sigma}_1} \int_{\underline{\theta}}^{\bar{\theta}} \theta \phi\left(\frac{\theta - \hat{\theta}_1^*}{\hat{\sigma}_1}\right) \left(1 - \Phi\left(\frac{\sigma_2(\hat{\theta}_2^* - \tilde{\theta}_2)}{\hat{\sigma}_2^2}\right)\right) d\theta + \frac{1}{\hat{\sigma}_1} \int_{\bar{\theta}}^{\infty} \theta \phi\left(\frac{\theta - \hat{\theta}_1^*}{\hat{\sigma}_1}\right) d\theta - T = 0 \quad (10)$$

$$v_2(\hat{\theta}_2^*, \hat{\theta}_1^*; \boldsymbol{\sigma}) \equiv \frac{1}{\hat{\sigma}_2} \int_{\underline{\theta}}^{\bar{\theta}} \theta \phi\left(\frac{\theta - \hat{\theta}_2^*}{\hat{\sigma}_2}\right) \left(1 - \Phi\left(\frac{\sigma_1(\hat{\theta}_1^* - \tilde{\theta}_1)}{\hat{\sigma}_1^2}\right)\right) d\theta + \frac{1}{\hat{\sigma}_2} \int_{\bar{\theta}}^{\infty} \theta \phi\left(\frac{\theta - \hat{\theta}_2^*}{\hat{\sigma}_2}\right) d\theta - T = 0 \quad (11)$$

where $\hat{\theta}_i^* = \frac{\sigma_i^2 \mu_\theta + \sigma_\theta^2 x_i^*}{\sigma_i^2 + \sigma_\theta^2}$, $\hat{\sigma}_i = \sqrt{\frac{\sigma_i^2 \sigma_\theta^2}{\sigma_i^2 + \sigma_\theta^2}}$ and $\tilde{\theta}_i = \frac{\sigma_i^2 \mu_\theta + \sigma_\theta^2 \theta}{\sigma_i^2 + \sigma_\theta^2}$ for $i = 1, 2$. Note that both equations determine $\hat{\theta}_j^*$ in terms of $\hat{\theta}_i^*$. Without loss of generality we analyze the behavior of $\hat{\theta}_2^*$ as a function of $\hat{\theta}_1^*$ in the $(\hat{\theta}_1^*, \hat{\theta}_2^*)$ space and rewrite equations (10) and (11) as:

$$v_1(\hat{\theta}_1^*, w_1(\hat{\theta}_1^*; \boldsymbol{\sigma}); \boldsymbol{\sigma}) = 0 \quad (12)$$

$$v_2(w_2(\hat{\theta}_1^*; \boldsymbol{\sigma}), \hat{\theta}_1^*; \boldsymbol{\sigma}) = 0 \quad (13)$$

where $\hat{\theta}_2^* = w_i(\hat{\theta}_1^*; \boldsymbol{\sigma})$ for $\hat{\theta}_2^*$ as defined by the equation that characterized player i 's payoff function, for $i = 1, 2$. Then any $\hat{\theta}_1^*$ that solves simultaneously both equations defines an equilibrium threshold for player 1 and the associated threshold for player 2 is simply given by $\hat{\theta}_2^* = w_1(\hat{\theta}_1^*; \boldsymbol{\sigma})$.

Consider first equation (12). Define $\bar{\theta}_1^*$ as the unique solution to the following equation:

$$\int_{\bar{\theta}}^{\infty} \theta \frac{1}{\hat{\sigma}_1} \phi \left(\frac{\theta - \bar{\theta}_1^*}{\hat{\sigma}_1} \right) d\theta - T = 0$$

Similarly, denote by $\underline{\theta}_1^*$ the unique solution to the following equation:

$$\int_{\underline{\theta}}^{\infty} \theta \frac{1}{\hat{\sigma}_1} \phi \left(\frac{\theta - \underline{\theta}_1^*}{\hat{\sigma}_1} \right) d\theta - T = 0$$

The first of the above conditions corresponds to the situation when player 2 never takes the risky actions while the second condition corresponds to the situation where player 2 always chooses to take the risky action. Note that $-\infty < \underline{\theta}_1^* < \bar{\theta}_1^* < \infty$ and therefore it follows that $\hat{\theta}_1^*$ is finite (and $\hat{\theta}_1^* \in [\underline{\theta}_1^*, \bar{\theta}_1^*]$). Recall that by lemma 1, the LHS of (12) is increasing in $\hat{\theta}_1^*$ and decreasing in $\bar{\theta}_2^*$. It follows then that as $\hat{\theta}_1^* \rightarrow \underline{\theta}_1^*$, $\bar{\theta}_2^* \rightarrow -\infty$ and as $\hat{\theta}_1^* \rightarrow \bar{\theta}_1^*$, $\bar{\theta}_2^* \rightarrow \infty$. Therefore $w_1(\hat{\theta}_1^*; \sigma)$ has asymptotes at $\underline{\theta}_1^*$ and $\bar{\theta}_1^*$. Similarly define $\underline{\theta}_2^*$ and $\bar{\theta}_2^*$ for player two. By lemma 1 we conclude that $w_2(\hat{\theta}_1^*; \sigma)$ is bounded above by $\bar{\theta}_2^*$ and below by $\underline{\theta}_2^*$. Finally, let $\theta_{\min}^* = \min\{\underline{\theta}_1^*, \underline{\theta}_2^*\}$ and $\theta_{\max}^* = \max\{\bar{\theta}_1^*, \bar{\theta}_2^*\}$ so that θ_{\min}^* is the smallest and θ_{\max}^* is the largest threshold that can be rationalized.

Using the implicit function theorem we can find the derivative of $w_1(\hat{\theta}_1^*; \sigma)$ wrt $\hat{\theta}_1^*$:

$$\frac{dw_1(\hat{\theta}_1^*)}{d\hat{\theta}_1^*} = \frac{\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\hat{\sigma}_1} \phi \left(\frac{\theta - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) \frac{1}{\sigma_2} \phi \left(\frac{\sigma_2(\hat{\theta}_2^* - \bar{\theta}_2)}{\hat{\sigma}_2^2} \right) d\theta + \tilde{V}_1}{\frac{\sigma_{\theta}^2 + \sigma_2^2}{\sigma_{\theta}^2} \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\hat{\sigma}_1} \phi \left(\frac{\theta - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) \frac{1}{\sigma_2} \phi \left(\frac{\sigma_2(\hat{\theta}_2^* - \bar{\theta}_2)}{\hat{\sigma}_2^2} \right) d\theta} > 0$$

where

$$\begin{aligned} \tilde{V}_1 &= 1 - \Phi \left(\frac{\bar{\theta} - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) + \bar{\theta} \frac{1}{\hat{\sigma}_1} \phi \left(\frac{\bar{\theta} - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) \Phi \left(\frac{\sigma_2 \left(\hat{\theta}_2^* - \frac{\sigma_{\theta}^2 \bar{\theta} + \sigma_2^2 \mu_{\theta}}{\sigma_{\theta}^2 + \sigma_2^2} \right)}{\hat{\sigma}_2^2} \right) \\ &\quad + \underline{\theta} \frac{1}{\hat{\sigma}_1} \phi \left(\frac{\underline{\theta} - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) \left(1 - \Phi \left(\frac{\sigma_2 \left(\hat{\theta}_2^* - \frac{\sigma_{\theta}^2 \underline{\theta} + \sigma_2^2 \mu_{\theta}}{\sigma_{\theta}^2 + \sigma_2^2} \right)}{\hat{\sigma}_2^2} \right) \right) \\ &\quad + \int_{\underline{\theta}}^{\bar{\theta}} \frac{1}{\hat{\sigma}_1} \phi \left(\frac{\theta - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) \left(1 - \Phi \left(\frac{\sigma_2(\hat{\theta}_2^* - \bar{\theta}_2)}{\hat{\sigma}_2^2} \right) \right) d\theta > 0 \end{aligned}$$

is strictly positive (since both $\hat{\theta}_1^*$ and $\hat{\theta}_2^*$ are finite - see the discussion above).

Similarly, we calculate the derivative of $w_2(\widehat{\theta}_1^*; \sigma)$ wrt $\widehat{\theta}_1^*$:

$$\frac{dw_2(\widehat{\theta}_1^*)}{d\widehat{\theta}_1^*} = \frac{\frac{\sigma_\theta^2 + \sigma_1^2}{\sigma_\theta^2} \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_2} \phi\left(\frac{\theta - \widehat{\theta}_2^*}{\sigma_2}\right) \frac{1}{\sigma_1} \phi\left(\frac{\sigma_1(\widehat{\theta}_1^* - \tilde{\theta}_1)}{\widehat{\sigma}_1^2}\right) d\theta}{\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_2} \phi\left(\frac{\theta - \widehat{\theta}_2^*}{\sigma_2}\right) \frac{1}{\sigma_1} \phi\left(\frac{\sigma_1(\widehat{\theta}_1^* - \tilde{\theta}_1)}{\widehat{\sigma}_1^2}\right) d\theta + \tilde{V}_2} > 0$$

where \tilde{V}_2 a strictly positive constant and is defined in analogously to \tilde{V}_1 .

Note that a sufficient condition for uniqueness is

$$\frac{dw_1(\widehat{\theta}_1^*)}{d\widehat{\theta}_1^*} > \frac{dw_2(\widehat{\theta}_1^*)}{d\widehat{\theta}_1^*} > 0$$

This translates in the following inequality:

$$\frac{\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_1} \phi\left(\frac{\theta - \widehat{\theta}_1^*}{\sigma_1}\right) \frac{1}{\sigma_2} \phi\left(\frac{\sigma_2(\widehat{\theta}_2^* - \tilde{\theta}_2)}{\widehat{\sigma}_2^2}\right) d\theta + \tilde{V}_1}{\frac{\sigma_\theta^2 + \sigma_2^2}{\sigma_\theta^2} \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_1} \phi\left(\frac{\theta - \widehat{\theta}_1^*}{\sigma_1}\right) \frac{1}{\sigma_2} \phi\left(\frac{\sigma_2(\widehat{\theta}_2^* - \tilde{\theta}_2)}{\widehat{\sigma}_2^2}\right) d\theta} > \frac{\frac{\sigma_\theta^2 + \sigma_1^2}{\sigma_\theta^2} \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_2} \phi\left(\frac{\theta - \widehat{\theta}_2^*}{\sigma_2}\right) \frac{1}{\sigma_1} \phi\left(\frac{\sigma_1(\widehat{\theta}_1^* - \tilde{\theta}_1)}{\widehat{\sigma}_1^2}\right) d\theta}{\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_2} \phi\left(\frac{\theta - \widehat{\theta}_2^*}{\sigma_2}\right) \frac{1}{\sigma_1} \phi\left(\frac{\sigma_1(\widehat{\theta}_1^* - \tilde{\theta}_1)}{\widehat{\sigma}_1^2}\right) d\theta + \tilde{V}_2}$$

Doing some algebraic manipulations, we get that the expression above is equivalent to

$$\begin{aligned} & \frac{\tilde{V}_1}{\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_1} \phi\left(\frac{\theta - \widehat{\theta}_1^*}{\sigma_1}\right) \frac{1}{\sigma_2} \phi\left(\frac{\sigma_2(\widehat{\theta}_2^* - \tilde{\theta}_2)}{\widehat{\sigma}_2^2}\right) d\theta} + \\ & \frac{\tilde{V}_2}{\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_2} \phi\left(\frac{\theta - \widehat{\theta}_2^*}{\sigma_2}\right) \frac{1}{\sigma_1} \phi\left(\frac{\sigma_1(\widehat{\theta}_1^* - \tilde{\theta}_1)}{\widehat{\sigma}_1^2}\right) d\theta} + \\ & \frac{\tilde{V}_1 \tilde{V}_2}{\left(\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_1} \phi\left(\frac{\theta - \widehat{\theta}_1^*}{\sigma_1}\right) \frac{1}{\sigma_2} \phi\left(\frac{\sigma_2(\widehat{\theta}_2^* - \tilde{\theta}_2)}{\widehat{\sigma}_2^2}\right) d\theta\right) \left(\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_2} \phi\left(\frac{\theta - \widehat{\theta}_2^*}{\sigma_2}\right) \frac{1}{\sigma_1} \phi\left(\frac{\sigma_1(\widehat{\theta}_1^* - \tilde{\theta}_1)}{\widehat{\sigma}_1^2}\right) d\theta\right)} \\ & > \frac{\sigma_1^2}{\sigma_\theta^2} + \frac{\sigma_2^2}{\sigma_\theta^2} + \frac{\sigma_1^2 \sigma_2^2}{\sigma_\theta^4} \end{aligned}$$

A sufficient condition for this inequality to hold is to have:

$$\frac{\sigma_1^2}{\sigma_\theta^2} < \frac{\tilde{V}_1}{\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_1} \phi\left(\frac{\theta - \widehat{\theta}_1^*}{\sigma_1}\right) \frac{1}{\sigma_2} \phi\left(\frac{\sigma_2(\widehat{\theta}_2^* - \tilde{\theta}_2)}{\widehat{\sigma}_2^2}\right) d\theta}$$

and

$$\frac{\sigma_2^2}{\sigma_\theta^2} < \frac{\tilde{V}_2}{\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_2} \phi\left(\frac{\theta - \widehat{\theta}_2^*}{\sigma_2}\right) \frac{1}{\sigma_1} \phi\left(\frac{\sigma_1(\widehat{\theta}_1^* - \tilde{\theta}_1)}{\widehat{\sigma}_1^2}\right) d\theta}$$

Take the first expression, for player 1 (the result is analogous for player 2). We want to find a lower bound for the RHS, i.e. a lower bound for the numerator of the RHS and an upper bound for the denominator of the RHS.

We consider first the denominator $\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_1} \phi\left(\frac{\theta - \widehat{\theta}_1^*}{\sigma_1}\right) \frac{1}{\sigma_2} \phi\left(\frac{\sigma_2(\widehat{\theta}_2^* - \tilde{\theta}_2)}{\widehat{\sigma}_2^2}\right) d\theta$. We can rewrite this term

as

$$\begin{aligned}
& \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\hat{\sigma}_1} \phi \left(\frac{\theta - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) \frac{1}{\sigma_2} \phi \left(\frac{\sigma_2(\hat{\theta}_2^* - \tilde{\theta}_2)}{\hat{\sigma}_2^2} \right) d\theta \\
&= \frac{1}{\sqrt{\hat{\sigma}_1^2 + \sigma_2^2}} \phi \left(\frac{\hat{\theta}_1^* - \Omega}{\sqrt{\hat{\sigma}_1^2 + \sigma_2^2}} \right) \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sqrt{\frac{\hat{\sigma}_1^2 \sigma_2^2}{\hat{\sigma}_1^2 + \sigma_2^2}}} \phi \left(\frac{\theta - \frac{\hat{\sigma}_1^2 \Omega + \sigma_2^2 \hat{\theta}_1^*}{\hat{\sigma}_1^2 + \sigma_2^2}}{\sqrt{\frac{\hat{\sigma}_1^2 \sigma_2^2}{\hat{\sigma}_1^2 + \sigma_2^2}}} \right) d\theta \\
&\leq \frac{1}{\sqrt{\hat{\sigma}_1^2 + \sigma_2^2}} \frac{1}{\sqrt{2\pi}} \left[\frac{\theta_{\max}^* \left(\hat{\sigma}_1^2 + \sigma_2^2 + \hat{\sigma}_1^2 \frac{\sigma_2^2}{\sigma_\theta^2} \right) - \hat{\sigma}_1^2 \frac{\sigma_2^2}{\sigma_\theta^2} \mu_\theta}{\hat{\sigma}_1^2 + \sigma_2^2} + \frac{1}{\sqrt{2\pi}} \frac{\sqrt{\hat{\sigma}_1^2 \sigma_2^2}}{\hat{\sigma}_1^2 + \sigma_2^2} \right]
\end{aligned}$$

where $\Omega \equiv \frac{(\sigma_\theta^2 + \sigma_2^2)}{\sigma_\theta^2} \hat{\theta}_2^* - \frac{\sigma_2^2}{\sigma_\theta^2} \mu_\theta$.

We now look at the numerator \tilde{V}_1 . Note that

$$\begin{aligned}
& \int_{\underline{\theta}}^{\bar{\theta}} \frac{1}{\hat{\sigma}_1} \phi \left(\frac{\theta - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) \left(1 - \Phi \left(\frac{\sigma_2(\hat{\theta}_2^* - \tilde{\theta}_2)}{\hat{\sigma}_2^2} \right) \right) d\theta \\
&\geq \left[\Phi \left(\frac{\bar{\theta} - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) - \Phi \left(\frac{\underline{\theta} - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) \right] \left(1 - \Phi \left(\frac{\sigma_2(\hat{\theta}_2^* - \tilde{\theta}_2)}{\hat{\sigma}_2^2} \right) \right)
\end{aligned}$$

where $\tilde{\theta}_2 \equiv \frac{\sigma_\theta^2 \bar{\theta} + \sigma_2^2 \mu_\theta}{\sigma_\theta^2 + \sigma_2^2}$ and therefore

$$\begin{aligned}
\tilde{V}_1 &> 1 - \Phi \left(\frac{\underline{\theta} - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) - \left[\Phi \left(\frac{\bar{\theta} - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) - \Phi \left(\frac{\underline{\theta} - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) \right] \Phi \left(\frac{\sigma_2(\hat{\theta}_2^* - \tilde{\theta}_2)}{\hat{\sigma}_2^2} \right) \\
&\quad + \bar{\theta} \frac{1}{\hat{\sigma}_1} \phi \left(\frac{\bar{\theta} - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) \Phi \left(\frac{\sigma_2 \left(\hat{\theta}_2^* - \frac{\sigma_\theta^2 \bar{\theta} + \sigma_2^2 \mu_\theta}{\sigma_\theta^2 + \sigma_2^2} \right)}{\hat{\sigma}_2^2} \right) + \underline{\theta} \frac{1}{\hat{\sigma}_1} \phi \left(\frac{\underline{\theta} - \hat{\theta}_1^*}{\hat{\sigma}_1} \right) \left(1 - \Phi \left(\frac{\sigma_2 \left(\hat{\theta}_2^* - \frac{\sigma_\theta^2 \underline{\theta} + \sigma_2^2 \mu_\theta}{\sigma_\theta^2 + \sigma_2^2} \right)}{\hat{\sigma}_2^2} \right) \right)
\end{aligned}$$

We can bound the last two terms of the RHS of the above expression by

$$\frac{1}{\hat{\sigma}_1} \phi \left(\frac{\bar{\theta} - \theta_{\min}^*}{\hat{\sigma}_1} \right) \left[\bar{\theta} \Phi \left(\frac{\sigma_2 \left(\theta_{\min}^* - \frac{\sigma_\theta^2 \bar{\theta} + \sigma_2^2 \mu_\theta}{\sigma_\theta^2 + \sigma_2^2} \right)}{\hat{\sigma}_2^2} \right) + \underline{\theta} \left(1 - \Phi \left(\frac{\sigma_2 \left(\theta_{\max}^* - \frac{\sigma_\theta^2 \underline{\theta} + \sigma_2^2 \mu_\theta}{\sigma_\theta^2 + \sigma_2^2} \right)}{\hat{\sigma}_2^2} \right) \right) \right]$$

Therefore, sufficient conditions for uniqueness are:

$$\frac{\sigma_2^2}{\sigma_\theta^2} < \frac{1 - \Phi\left(\frac{\theta - \hat{\theta}_{\min}}{\hat{\sigma}_1}\right) - \left[\Phi\left(\frac{\bar{\theta} - \hat{\theta}_{\min}}{\hat{\sigma}_1}\right) - \Phi\left(\frac{\theta - \hat{\theta}_{\max}}{\hat{\sigma}_1}\right)\right] \Phi\left(\frac{\sigma_2(\hat{\theta}_{\max} - \tilde{\theta}_2)}{\hat{\sigma}_2^2}\right) + \frac{1}{\hat{\sigma}_1} \phi\left(\frac{\bar{\theta} - \theta^*_{\min}}{\hat{\sigma}_1}\right) \kappa_2}{\frac{1}{\sqrt{2\pi}} \left(\frac{1}{\sqrt{2\pi}} + \frac{\bar{\theta}\sigma_2^2\left(1 + \frac{\sigma_2^2}{\hat{\sigma}_1^2}\right) - \hat{\sigma}_1^2 \frac{\sigma_1^2}{\sigma_\theta^2}}{\hat{\sigma}_1^2 + \sigma_2^2} \right)}$$

$$\frac{\sigma_1^2}{\sigma_\theta^2} < \frac{1 - \Phi\left(\frac{\bar{\theta} - \theta}{\hat{\sigma}_2}\right) + \frac{1}{\hat{\sigma}_2} \phi\left(\frac{\bar{\theta} - \theta}{\hat{\sigma}_2}\right) \kappa_1}{\frac{1}{\sqrt{2\pi}} \left(\frac{1}{\sqrt{2\pi}} + \frac{\bar{\theta}\sigma_1^2\left(1 + \frac{\sigma_1^2}{\hat{\sigma}_2^2}\right) - \hat{\sigma}_2^2 \frac{\sigma_2^2}{\sigma_\theta^2}}{\hat{\sigma}_2^2 + \sigma_1^2} \right)}$$

$$\text{Where } \kappa_i := \bar{\theta} \Phi\left(\frac{\sigma_2\left(\theta^*_{\min} - \frac{\sigma_\theta^2 \bar{\theta} + \sigma_2^2 \mu_\theta}{\sigma_\theta^2 + \sigma_2^2}\right)}{\hat{\sigma}_2^2}\right) + \theta \left(1 - \Phi\left(\frac{\sigma_2\left(\theta^*_{\max} - \frac{\sigma_\theta^2 \theta + \sigma_2^2 \mu_\theta}{\sigma_\theta^2 + \sigma_2^2}\right)}{\hat{\sigma}_2^2}\right)\right)$$

If such conditions hold, $0 < \frac{dw_2(\hat{\theta}_1^*; \boldsymbol{\sigma})}{d\hat{\theta}_1^*} < \frac{dw_1(\hat{\theta}_1^*; \boldsymbol{\sigma})}{d\hat{\theta}_1^*} \forall \hat{\theta}_1^* \in [\theta_1^*, \bar{\theta}_1^*]$. This means that the least and greatest Bayesian Nash equilibria of the game, as described by our Corollary in the first section of the appendix, coincide. Therefore, there is a unique equilibrium in thresholds strategies. This proves the first part of the theorem.

The proof for the second part of the theorem follows directly from the proof of the above result. Namely, recall that to prove uniqueness we have to find conditions under which the functions $w_1(\hat{\theta}_1^*)$ and $w_2(\hat{\theta}_1^*)$ are such that

$$\frac{dw_1(\hat{\theta}_1^*)}{d\hat{\theta}_1^*} > \frac{dw_2(\hat{\theta}_1^*)}{d\hat{\theta}_1^*}$$

Note that as $\sigma_\theta \rightarrow \infty$ we have

$$\lim_{\sigma_\theta \rightarrow \infty} \frac{\sigma_\theta^2 + \sigma_1^2}{\sigma_\theta^2} \rightarrow 1 \text{ and } \lim_{\sigma_\theta \rightarrow \infty} \frac{\sigma_\theta^2 + \sigma_2^2}{\sigma_\theta^2} \rightarrow 1$$

Therefore

$$\begin{aligned} \frac{dw_1(\hat{\theta}_1^*)}{d\hat{\theta}_1^*} &= \frac{\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\hat{\sigma}_1} \phi\left(\frac{\theta - \hat{\theta}_1^*}{\hat{\sigma}_1}\right) \frac{1}{\sigma_2} \phi\left(\frac{\sigma_2(\hat{\theta}_2^* - \tilde{\theta}_2)}{\hat{\sigma}_2^2}\right) d\theta + \tilde{V}_1}{\frac{\sigma_\theta^2 + \sigma_2^2}{\sigma_\theta^2} \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\hat{\sigma}_1} \phi\left(\frac{\theta - \hat{\theta}_1^*}{\hat{\sigma}_1}\right) \frac{1}{\sigma_2} \phi\left(\frac{\sigma_2(\hat{\theta}_2^* - \tilde{\theta}_2)}{\hat{\sigma}_2^2}\right) d\theta} \\ &\rightarrow \frac{\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_1} \phi\left(\frac{\theta - x_1^*}{\sigma_1}\right) \frac{1}{\sigma_2} \phi\left(\frac{x_2^* - \theta}{\sigma_2}\right) d\theta + \tilde{V}_1}{\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\sigma_1} \phi\left(\frac{\theta - x_1^*}{\sigma_1}\right) \frac{1}{\sigma_2} \phi\left(\frac{x_2^* - \theta}{\sigma_2}\right) d\theta} > 1 \end{aligned}$$

and

$$\begin{aligned} \frac{dw_2(\hat{\theta}_1^*)}{d\hat{\theta}_1^*} &= \frac{\frac{\sigma_\theta^2 + \sigma_1^2}{\sigma_\theta^2} \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\hat{\sigma}_2} \phi\left(\frac{\theta - \hat{\theta}_2^*}{\hat{\sigma}_2}\right) \frac{1}{\sigma_1} \phi\left(\frac{\sigma_1(\hat{\theta}_1^* - \tilde{\theta}_1)}{\hat{\sigma}_1^2}\right) d\theta}{\tilde{V}_2 + \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\hat{\sigma}_2} \phi\left(\frac{\theta - \hat{\theta}_2^*}{\hat{\sigma}_2}\right) \frac{1}{\sigma_1} \phi\left(\frac{\sigma_1(\hat{\theta}_1^* - \tilde{\theta}_1)}{\hat{\sigma}_1^2}\right) d\theta} \\ &\rightarrow \frac{\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\hat{\sigma}_2} \phi\left(\frac{\theta - x_2^*}{\hat{\sigma}_2}\right) \frac{1}{\sigma_1} \phi\left(\frac{x_1^* - \theta}{\sigma_1}\right) d\theta}{\tilde{V}_2 + \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\hat{\sigma}_2} \phi\left(\frac{\theta - x_2^*}{\hat{\sigma}_2}\right) \frac{1}{\sigma_1} \phi\left(\frac{x_1^* - \theta}{\sigma_1}\right) d\theta} < 1 \end{aligned}$$

since, as we argued in the proof of the first part of the theorem, \tilde{V}_1 and \tilde{V}_2 are strictly positive. By continuity of the above expressions we conclude that for any σ_1 and σ_2 there exists $\bar{\sigma}_\theta(\sigma_1, \sigma_2)$ such that if $\sigma_\theta > \bar{\sigma}_\theta(\sigma_1, \sigma_2)$ we have a unique equilibrium in the coordination game.

The above bound depends on the information acquisition choices made by players. However, it is easy to show existence of a uniform bound. To do so, recall first that

$$\begin{aligned} \tilde{V}_1 &= 1 - \Phi\left(\frac{\bar{\theta} - \hat{\theta}_1^*}{\hat{\sigma}_1}\right) + \bar{\theta} \frac{1}{\hat{\sigma}_1} \phi\left(\frac{\bar{\theta} - \hat{\theta}_1^*}{\hat{\sigma}_1}\right) \Phi\left(\frac{\sigma_2\left(\hat{\theta}_2^* - \frac{\sigma_\theta^2 \bar{\theta} + \sigma_2^2 \mu_\theta}{\sigma_\theta^2 + \sigma_2^2}\right)}{\hat{\sigma}_2^2}\right) \\ &\quad + \frac{\theta}{\hat{\sigma}_1} \phi\left(\frac{\theta - \hat{\theta}_1^*}{\hat{\sigma}_1}\right) \left(1 - \Phi\left(\frac{\sigma_2\left(\hat{\theta}_2^* - \frac{\sigma_\theta^2 \theta + \sigma_2^2 \mu_\theta}{\sigma_\theta^2 + \sigma_2^2}\right)}{\hat{\sigma}_2^2}\right)\right) \\ &\quad + \int_{\underline{\theta}}^{\bar{\theta}} \frac{1}{\hat{\sigma}_1} \phi\left(\frac{\theta - \hat{\theta}_1^*}{\hat{\sigma}_1}\right) \left(1 - \Phi\left(\frac{\sigma_2(\hat{\theta}_2^* - \tilde{\theta}_2)}{\hat{\sigma}_2^2}\right)\right) d\theta > 0 \end{aligned}$$

and in particular

$$\tilde{V}_1 > 1 - \Phi\left(\frac{\bar{\theta} - \hat{\theta}_1^*}{\hat{\sigma}_1}\right) > 0$$

Recall that $\hat{\sigma}_1 = \sqrt{\frac{\sigma_1^2 \sigma_\theta^2}{\sigma_1^2 + \sigma_\theta^2}}$ and that $\sigma_1^2 \leq \sigma_0^2$ where σ_0^2 is the precision of a private signal if player 1 did not acquire any information. Let $\sigma_\theta^{2,B}$ be an arbitrary lower bound on σ_θ^2 such that $\sigma_\theta^{2,B} > 0$. If at $\sigma_\theta^2 = \sigma_\theta^{2,B}$ we have $\frac{dw_1(\hat{\theta}_1^*)}{d\hat{\theta}_1^*} > \frac{dw_2(\hat{\theta}_1^*)}{d\hat{\theta}_1^*}$ then we are done. Otherwise, we have to show that there exists a bound on σ_θ^2 that is higher than $\sigma_\theta^{2,B}$ for which $\frac{dw_1(\hat{\theta}_1^*)}{d\hat{\theta}_1^*} > \frac{dw_2(\hat{\theta}_1^*)}{d\hat{\theta}_1^*}$ independent of values of σ_1 and σ_2 .

To do so we start by finding uniform bounds on $\hat{\theta}_1^*$. Note that, for any σ_1 and σ_θ , the lowest threshold that player 1 can possibly choose (which we denote by $\underline{\theta}_1^*(\sigma_\theta, \sigma_1)$) is determined by equation

$$\int_{\underline{\theta}}^{\infty} \frac{1}{\hat{\sigma}_1^2} \theta \phi\left(\frac{\theta - \underline{\theta}_1^*}{\hat{\sigma}_1}\right) d\theta = T$$

which corresponds to the situation in which the other player always chooses to take the risky action

(and where we suppressed the dependence of $\underline{\theta}_1^*$ on $(\sigma_\theta, \sigma_1)$). This can be written as

$$\underline{\theta}_1^* \left(1 - \Phi \left(\frac{\theta - \underline{\theta}_1^*}{\widehat{\sigma}_1} \right) \right) + \widehat{\sigma}_1 \phi \left(\frac{\theta - \underline{\theta}_1^*}{\widehat{\sigma}_1} \right) = T \quad (14)$$

By the implicit function theorem we have

$$\frac{\partial \underline{\theta}_1^*}{\partial \widehat{\sigma}_1} = - \frac{\underline{\theta}_1^* \phi \left(\frac{\theta - \underline{\theta}_1^*}{\widehat{\sigma}_1} \right) \frac{\theta - \underline{\theta}_1^*}{\widehat{\sigma}_1^2} + \phi \left(\frac{\theta - \underline{\theta}_1^*}{\widehat{\sigma}_1} \right) + \frac{(\theta - \underline{\theta}_1^*)^2}{\widehat{\sigma}_1^2} \phi \left(\frac{\theta - \underline{\theta}_1^*}{\widehat{\sigma}_1} \right)}{\left(1 - \Phi \left(\frac{\theta - \underline{\theta}_1^*}{\widehat{\sigma}_1} \right) \right) + \frac{\underline{\theta}_1^*}{\widehat{\sigma}_1} \phi \left(\frac{\theta - \underline{\theta}_1^*}{\widehat{\sigma}_1} \right) + \frac{\theta - \underline{\theta}_1^*}{\widehat{\sigma}_1} \phi \left(\frac{\theta - \underline{\theta}_1^*}{\widehat{\sigma}_1} \right)} < 0$$

So as $\widehat{\sigma}_1$ increases $\underline{\theta}_1^*$ decreases, which implies that an increase in σ_θ^2 decreases $\underline{\theta}_1^*$ while an increase in σ_1 increases $\underline{\theta}_1^*$. In the same way we can show that the highest threshold that player 1 can possibly choose (denoted by $\bar{\theta}_1^*$) is also decreasing in $\widehat{\sigma}_1$. Therefore, $\underline{\theta}_1^*$ is minimized at $\sigma_1 = \sigma_0$ and $\sigma_\theta \rightarrow \infty$. This implies that

$$\tilde{V}_1 > 1 - \Phi \left(\frac{\bar{\theta} - \widehat{\theta}_1^*}{\widehat{\sigma}_1} \right) > K_1$$

where

$$K_1 \equiv 1 - \Phi \left(\frac{\bar{\theta} - \lim_{\sigma_\theta \rightarrow \infty} \underline{\theta}_1^*(\sigma_\theta, \sigma_0)}{\widehat{\sigma}_1} \right)$$

This establishes a bound on \tilde{V}_1 that is independent of σ_θ, σ_0 . Note that by symmetry this is also a lower bound on \tilde{V}_2 .

We now show that we can also bound other term appearing in the expression for the derivative uniformly. But

$$\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\widehat{\sigma}_2} \phi \left(\frac{\theta - \widehat{\theta}_2^*}{\widehat{\sigma}_2} \right) \frac{1}{\sigma_1} \phi \left(\frac{\sigma_1(\widehat{\theta}_1^* - \tilde{\theta}_1)}{\widehat{\sigma}_1^2} \right) d\theta < \frac{1}{2\pi} [\bar{\theta} - \underline{\theta}] \frac{1}{\widehat{\sigma}_2} \frac{1}{\sigma_1}$$

Below, when we discuss the first stage of the game we show that the benefit from acquiring information tends to zero as $\sigma_i \rightarrow 0$ and therefore, given our assumptions on the cost function, i.e. $C'(\sigma_i) > 0$ and $\lim_{\sigma_i \rightarrow \infty} C'(\sigma_i) \rightarrow \infty$, there is a bound on the precision choice, call it σ_i^{\min} , such that player i will never choose to acquire a lower standard deviation than σ_i^{\min} . Therefore,

$$\begin{aligned} \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\widehat{\sigma}_2} \phi \left(\frac{\theta - \widehat{\theta}_2^*}{\widehat{\sigma}_2} \right) \frac{1}{\sigma_1} \phi \left(\frac{\sigma_1(\widehat{\theta}_1^* - \tilde{\theta}_1)}{\widehat{\sigma}_1^2} \right) &< \frac{1}{2\pi} [\bar{\theta} - \underline{\theta}] \frac{1}{\widehat{\sigma}_2} \frac{1}{\sigma_1} \\ &\leq \frac{1}{\widehat{\sigma}_2} \frac{1}{\sigma_1^{\min}} \frac{1}{2\pi} [\bar{\theta} - \underline{\theta}] \end{aligned}$$

Finally, note that $\widehat{\sigma}_2$ is increasing in σ_θ and σ_2 and so $\widehat{\sigma}_2$ is minimized at $\sigma_\theta = \sigma_\theta^B$ (our exogenous lower bound on σ_θ) and $\sigma_2 = \sigma_2^{\min}$ and denote by $\widehat{\sigma}_2^{\min}$ the posterior standard deviation of θ for

player 2 when $\sigma_\theta = \sigma_\theta^B$ and $\sigma_2 = \sigma_2^{\min}$. Then

$$\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\widehat{\sigma}_2} \phi\left(\frac{\theta - \widehat{\theta}_2^*}{\widehat{\sigma}_2}\right) \frac{1}{\sigma_1} \phi\left(\frac{\sigma_1(\widehat{\theta}_1^* - \widetilde{\theta}_1)}{\widehat{\sigma}_1^2}\right) \leq K_2$$

where

$$K_2 \equiv \frac{1}{\widehat{\sigma}_2^{\min}} \frac{1}{\sigma_1^{\min}} \frac{1}{2\pi} [\bar{\theta} - \underline{\theta}]$$

Therefore, we have

$$\begin{aligned} \frac{dw_1(\widehat{\theta}_1^*)}{d\widehat{\theta}_1^*} &= \frac{\int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\widehat{\sigma}_1} \phi\left(\frac{\theta - \widehat{\theta}_1^*}{\widehat{\sigma}_1}\right) \frac{1}{\sigma_2} \phi\left(\frac{\sigma_2(\widehat{\theta}_2^* - \widetilde{\theta}_2)}{\widehat{\sigma}_2^2}\right) d\theta + \widetilde{V}_1}{\frac{\sigma_\theta^2 + \sigma_2^2}{\sigma_\theta^2} \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\widehat{\sigma}_1} \phi\left(\frac{\theta - \widehat{\theta}_1^*}{\widehat{\sigma}_1}\right) \frac{1}{\sigma_2} \phi\left(\frac{\sigma_2(\widehat{\theta}_2^* - \widetilde{\theta}_2)}{\widehat{\sigma}_2^2}\right) d\theta} \\ &= \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_2^2} + \frac{\widetilde{V}_1}{\frac{\sigma_\theta^2 + \sigma_2^2}{\sigma_\theta^2} \int_{\underline{\theta}}^{\bar{\theta}} \theta \frac{1}{\widehat{\sigma}_1} \phi\left(\frac{\theta - \widehat{\theta}_1^*}{\widehat{\sigma}_1}\right) \frac{1}{\sigma_2} \phi\left(\frac{\sigma_2(\widehat{\theta}_2^* - \widetilde{\theta}_2)}{\widehat{\sigma}_2^2}\right) d\theta} \\ &\geq \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_2^2} + \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_2^2} \frac{K_1}{K_2} \\ &\geq \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_0^2} + \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_0^2} \frac{K_1}{K_2} \end{aligned}$$

where we used the fact that $\sigma_2^2 \leq \sigma_0^2$. Note that as $\sigma_\theta^2 \rightarrow \infty$ we have $\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_0^2} + \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_0^2} \frac{K_1}{K_2} \rightarrow 1 + \frac{K_1}{K_2} > 1$ and therefore there exists a bound on σ_θ^2 , call it $\bar{\sigma}_\theta^{2,P1}$ such that if $\sigma_\theta^2 > \bar{\sigma}_\theta^{2,P1}$ then $\frac{dw_1(\widehat{\theta}_1^*)}{d\widehat{\theta}_1^*} > 1$ irrespective of σ_1 and σ_2 .⁶³

Following the same steps as above we can show that there exists a bound on σ_θ^2 , which we denote by $\bar{\sigma}_\theta^{2,P2}$, such that if $\sigma_\theta^2 > \bar{\sigma}_\theta^{2,P2}$ then $\frac{dw_2(\widehat{\theta}_1^*)}{d\widehat{\theta}_1^*} < 1$ and $\bar{\sigma}_\theta^{2,P2}$ is independent of σ_1 and σ_2 .

Setting $\bar{\sigma}_\theta = \max\{\bar{\sigma}_\theta^{2,P1}, \bar{\sigma}_\theta^{2,P2}\}$ proves the second part of the theorem. ■

Lemma 2 *Suppose that $\sigma_i \rightarrow 0$, $\sigma_j \rightarrow 0$ and $\frac{\sigma_i}{\sigma_j} \rightarrow c$ where $c \in \overline{\mathbb{R}}_+$. If the above game has a unique equilibrium then this equilibrium converges to the risk-dominant equilibrium of the complete information game, i.e. $x_i^* \rightarrow 2T$ and $x_j^* \rightarrow 2T$.*

Proof. Player i 's expected payoff of taking the risky action is given by:

$$\int_{\underline{\theta}}^{\bar{\theta}} \theta \left(1 - \Phi\left(\frac{x_j^* - \theta}{\tau_{2j}^{-1/2}}\right)\right) (\tau_i + \tau_2)^{1/2} \phi\left(\frac{\theta - \frac{\tau_i x_i^* + \tau_\theta \mu_\theta}{\tau_i + \tau_\theta}}{(\tau_i + \tau_\theta)^{-1/2}}\right) d\theta + \int_{\bar{\theta}}^{\infty} \theta (\tau_i + \tau_2)^{1/2} \phi\left(\frac{\theta - \frac{\tau_i x_i^* + \tau_\theta \mu_\theta}{\tau_i + \tau_\theta}}{(\tau_i + \tau_\theta)^{-1/2}}\right) d\theta - T = 0$$

⁶³We use the superscript $P1$ to emphasize the fact that this restriction follows from equilibrium condition of player 1.

where $\tau_i = \sigma_i^{-2}$, $i = 1, 2$. We perform the following substitution in the above integrals

$$z = (\tau_i + \tau_\theta)^{1/2} \left[\theta - \frac{\tau_i x_i^* + \tau_\theta \mu_\theta}{\tau_i + \tau_\theta} \right]$$

Then, the indifference condition of player i becomes

$$\begin{aligned} & \int_{\underline{L}}^{\bar{L}} \left[\frac{z}{(\tau_i + \tau_\theta)^{1/2}} + \frac{\tau_i x_i^* + \tau_\theta \mu_\theta}{\tau_i + \tau_\theta} \right] \left(1 - \Phi \left(\frac{x_j^* - \frac{z}{(\tau_i + \tau_\theta)^{1/2}} - \frac{\tau_i x_i^* + \tau_\theta \mu_\theta}{\tau_i + \tau_\theta}}{\tau_{2j}^{-1/2}} \right) \right) \phi(z) dz \\ & + \int_{\underline{L}}^{\infty} \left[\frac{z}{(\tau_i + \tau_\theta)^{1/2}} + \frac{\tau_i x_i^* + \tau_\theta \mu_\theta}{\tau_i + \tau_\theta} \right] \phi(z) dz - T = 0 \end{aligned}$$

where

$$\underline{L} = \frac{\theta - \frac{\tau_i x_i^* + \tau_\theta \mu_\theta}{\tau_i + \tau_\theta}}{(\tau_i + \tau_\theta)^{-1/2}} \quad \text{and} \quad \bar{L} = \frac{\bar{\theta} - \frac{\tau_i x_i^* + \tau_\theta \mu_\theta}{\tau_i + \tau_\theta}}{(\tau_i + \tau_\theta)^{-1/2}}$$

We now take the limit as τ_1 and τ_2 tend to ∞ and $\tau_1/\tau_2 \rightarrow c \in \mathbb{R}$. We suppose that $\lim_{\tau_i \rightarrow \infty} x_i^* < \bar{\theta}$ in which case, the indifference conditions become

$$\begin{aligned} & x_i^{*,\text{lim}} \left(1 - \Phi \left(\lim_{\substack{\tau_i, \tau_j \rightarrow \infty \\ \tau_i/\tau_j \rightarrow c}} \frac{x_j^* - x_i^*}{\tau_j^{-1/2}} \right) \right) - T = 0 \\ & x_j^{*,\text{lim}} \left(1 - \Phi \left(\lim_{\substack{\tau_i, \tau_j \rightarrow \infty \\ \tau_i/\tau_j \rightarrow c}} \left(\frac{\tau_i}{\tau_j} \right)^{1/2} \frac{x_i^* - x_j^*}{\tau_j^{-1/2}} \right) \right) - T = 0 \end{aligned}$$

where $x_i^{*,\text{lim}}$ and $x_j^{*,\text{lim}}$ denote the thresholds in the limit. It is easy to see that in the limit we must have $x_i^{*,\text{lim}} = x_j^{*,\text{lim}}$.⁶⁴ But if $x_i^{*,\text{lim}} = x_j^{*,\text{lim}}$ then the above equations imply that $\tau_j^{1/2} (x_j^* - x_i^*) \rightarrow 0$. It follows that $x_i^{*,\text{lim}} = x_j^{*,\text{lim}} = 2T$.

The above argument was made under assumption that $x_i^{*,\text{lim}} < \bar{\theta}$ and $x_j^{*,\text{lim}} < \bar{\theta}$. It is straightforward to see that it cannot be the case that $x_i^{*,\text{lim}} > \bar{\theta}$ and $x_j^{*,\text{lim}} > \bar{\theta}$. Thus, suppose that $x_i^{*,\text{lim}} = x_j^{*,\text{lim}} = \bar{\theta}$. In that case the indifference condition in the limit as $\tau_i, \tau_j \rightarrow \infty$ and $\tau_i/\tau_j \rightarrow c \in \mathbb{R}$ converges to

$$\int_{-\infty}^L \bar{\theta} \left(1 - \Phi \left(\lim_{\substack{\tau_i, \tau_j \rightarrow \infty \\ \tau_i/\tau_j \rightarrow c}} \frac{x_j^* - x_i^*}{\tau_j^{-1/2}} \right) \right) \phi(z) dz + \int_L^{\infty} \bar{\theta} \phi(z) dz - T = 0$$

⁶⁴If for example $x_i^{*,\text{lim}} > x_j^{*,\text{lim}}$ then the above equations imply that $x_i^{*,\text{lim}} = T$ while $x_j^{*,\text{lim}} = \infty$ which cannot be the case as a player in the limit will never use a threshold larger than $\bar{\theta}$

where

$$L = \lim_{\substack{\tau_i, \tau_j \rightarrow \infty \\ \tau_i/\tau_j \rightarrow c}} \frac{\bar{\theta} - \frac{\tau_i x_i^* + \tau_\theta \mu_\theta}{\tau_i + \tau_\theta}}{(\tau_i + \tau_\theta)^{-1/2}}$$

The LHS of the resulting indifference equation is decreasing in L . Thus, if we can show that at $L = \infty$ the LHS is strictly larger than 0 we would arrive at a contradiction. But note that as $L = \infty$ then the indifference equations are given by

$$\int_{-\infty}^{\infty} \bar{\theta} \left(1 - \Phi \left(\lim_{\substack{\tau_i, \tau_j \rightarrow \infty \\ \tau_i/\tau_j \rightarrow c}} \frac{x_j^* - x_i^*}{\tau_j^{-1/2}} \right) \right) \phi(z) dz = T \quad \text{and} \quad \int_{-\infty}^{\infty} \bar{\theta} \left(1 - \Phi \left(\lim_{\substack{\tau_i, \tau_j \rightarrow \infty \\ \tau_i/\tau_j \rightarrow c}} \left(\frac{\tau_i}{\tau_j} \right)^{1/2} \frac{x_i^* - x_j^*}{\tau_j^{-1/2}} \right) \right) \phi(z) dz = T$$

These two equations can be satisfied simultaneously if and only if $\tau_j^{-1/2} (x_j^* - x_i^*) \rightarrow 0$ which implies that $\bar{\theta} = 2T$. But this is a contradiction since $\bar{\theta} > 2T$. So $x_i^{*, \text{lim}} = x_j^{*, \text{lim}} < \bar{\theta}$.

It is easy to verify that $2T$ corresponds to a threshold such that if $\theta \geq 2T$ then taking the risky action is a risk-dominant action and if $\theta < 2T$ not taking the safe action is a risk-dominant function. ■

C.1.3 Model with costly information acquisition

Note that the analysis of the model with heterogenous information described above corresponds to the equilibrium of the second stage of the model with costly information acquisition, conditional on precision choices made by players in the first stage. Thus, what remains is to solve for the information acquisition stage.

Theorem (Existence) *There exists a symmetric pure-strategy Bayesian Nash Equilibrium of the game with information acquisition.*

We first prove two simple claims and two corollaries that will make the proof of existence straightforward.

Claim 1 *$B_i(\sigma_i; \sigma_j)$ is a decreasing function of σ_i , that is $B_{i1}(\sigma_i; \sigma_j) \leq 0$.*

Proof. Since $u(\theta, a_i)$ has the single crossing property in (θ, a_i) , and that the signal x_i and the unknown parameter θ are affiliated, the claim then follows from Theorem 1 in Persico (2000). ■

Claim 2 *The marginal benefit of increasing the precision of player i converges to zero as the signal noise for player i vanishes, i.e. $\lim_{\sigma_i \rightarrow 0} \frac{\partial}{\partial \sigma_i} B_i(\sigma_i; \sigma_j) = 0$.*

This proof is very lengthy and can be obtained from the authors by request. It requires to show that in the limit $\frac{\partial}{\partial \sigma_i} x_i^*(\sigma_i, \sigma_j)$ is bounded and to verify that all the integrals in the expression for the marginal benefit converge to zero.

From the above results we have the following immediate corollaries:

Corollary 1 *The best response functions for both players are well defined.*

Proof. Since the cost function is strictly decreasing in σ_i and tends to infinite as $\sigma_i \rightarrow 0$, and since $B_i(\sigma_i, \sigma_j)$ is positive and stays bounded for each σ_j , we know that for each σ_j there is a unique choice of σ_i , holding beliefs of both players constant. ■

Corollary 2 *In any equilibrium of the game both players choose to acquire information (increase the precision of their signals).*

Proof. This follows from the fact that the marginal cost of acquiring information is continuous and equal zero at $\sigma_i = \sigma_0$, together with the fact that the marginal benefit of lowering σ_i is strictly positive for $\sigma_i > 0$.⁶⁵ ■

Existence. Suppose that player j believes that whenever he chooses a precision σ_j , player i will make the same choice. Holding player j 's beliefs fixed, we showed above that the best response function

$$\sigma_i^*(\sigma_j) = \max_{\sigma_i \in [0, \sigma_0]} U_i(\sigma_i, \sigma_j)$$

is well defined. Since $U_i(\sigma_i, \sigma_j)$ is continuous in σ_i and σ_j , by the Theorem of the Maximum we conclude that $\sigma_i^*(\sigma_j)$ is a continuous function. $\sigma_i^*(\cdot)$ is also a self-map: $\sigma_i^* : [0, \sigma_0] \rightarrow [0, \sigma_0]$. Hence, by Brower's Fixed Point Theorem, $\sigma_i^*(\cdot)$ has a fixed point. This implies that there exists a σ_j such that if player j believes that player i chooses $\sigma_i = \sigma_j$ player i will find it optimal to choose such a σ_i , that is $\sigma_i^*(\sigma_j) = \sigma_j$. ■

D Proofs of the model with biases in belief formation - FOR ON-LINE PUBLICATION

In this Section we prove Proposition 3 from the main text. We first show that the equilibrium exists and is unique in the limit as $\sigma \rightarrow 0$. Then, we show that there exists $\bar{\sigma} > 0$ such that for all $\sigma < \bar{\sigma}$ the model with biased beliefs has a unique equilibrium. We find it more convenient to work with precisions of signals and the prior, denoted by τ_x and τ_θ , respectively, rather than standard deviations.⁶⁶

We will make use of two results:

Lemma 7 *Consider a sequence $\{c_n\}_{n=1}^\infty \subset [a, b]$. If every convergent subsequence of c_n has the same limit L , then c_n is convergent sequence with a limit L .*

Note that this result allows us to focus on convergent subsequences in our analysis of limiting behavior of thresholds and then, if we confirm the hypothesis of the above Lemma, we know this

⁶⁵Note that above we established that $B_{i1} \leq 0$ and $\lim_{\sigma_i \rightarrow 0} B_{i1} = 0$. It can be shown that $B_{i1}(s_i, \sigma_j) \neq B_{i1}(s'_i, \sigma_j) \forall s_i \neq s'_i$. It follows then that $B_{i1}(\sigma_i, \sigma_j) < 0$ whenever $\sigma_i > 0$. That is, decreasing σ_i (increasing the precision) strictly increases the gross payoff for player i .

⁶⁶Recall that precision of a random variable distributed according to $N(\mu, \sigma^2)$ is defined as $\tau = \sigma^{-2}$.

limiting behavior is passed onto the sequence itself.

Lemma 8 *Let $F : \Omega \rightarrow \mathbb{R}^n$ where Ω is a rectangular region of \mathbb{R}^n . Let $J(x)$ be the Jacobian of the mapping F . If $J(x)$ is a diagonally dominant matrix with strictly positive diagonal then F is globally univalent on Ω .⁶⁷*

Proof. See Parthasarathy (1983). ■

Suppose that players use thresholds $\{x^*(\alpha_l)\}_{l=1}^N$ where $x^*(\alpha_l)$ is the threshold used by a player with bias α_l . Given the vector of thresholds $\{x^*(\alpha_k)\}_{k=1}^N$ the indifference condition for a player with bias α_k is given

$$V(x^*(\alpha_k), \{x^*(\alpha_l)\}_{l=1}^N) = 0$$

where

$$V(x^*(\alpha_k), \{x^*(\alpha_l)\}_{l=1}^N) = \sum_{l=1}^N g(\alpha_l) \int_{\underline{\theta}}^{\bar{\theta}} \theta \left[1 - \Phi \left(\frac{x^*(\alpha_l) - \theta}{\tau_x^{-1/2}} - \alpha_k \right) \right] f(\theta|x^*(\alpha_k)) d\theta + \int_{\bar{\theta}}^{\infty} \theta f(\theta|x^*(\alpha_k)) d-T$$

and

$$f(\theta|x^*(\alpha_k)) = (\tau_x + \tau_\theta)^{1/2} \phi \left(\frac{\theta - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}} \right)$$

Our goal is to show that the system of equations defined by

$$\begin{aligned} V(x^*(\alpha_1), \{x^*(\alpha_l)\}_{l=1}^N) &= 0 \\ &\vdots \\ V(x^*(\alpha_N), \{x^*(\alpha_l)\}_{l=1}^N) &= 0 \end{aligned}$$

has a unique solution. The proof boils down to establishing that the mapping $\mathbf{V} : \mathbb{R}^N \rightarrow \mathbb{R}^N$ defined by the LHS of the above system of equations is univalent, which, by the Gale-Nikaido Theorem, implies that there exists a unique vector of thresholds that satisfies the above system of indifference conditions. As Lemma 8 indicates, it is enough to show that the Jacobian of \mathbf{V} is diagonally dominant.

D.1 Equilibrium in the limit as $\tau_x \rightarrow \infty$

We first establish uniqueness of equilibrium in the limit as the precision of private signals tends to infinity. Throughout this section it is convenient to make the following substitution in the players'

⁶⁷ A square matrix A is called strictly diagonally dominant if $|A_{ii}| > \sum_{j \neq i} |A_{ij}|$ for each row i .

indifference conditions:

$$z = \frac{\theta - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}}$$

The payoff indifference condition becomes

$$\begin{aligned} \sum_{l=1}^N g(\alpha_l) \int_{L(\underline{\theta}, x^*(\alpha_k))}^{L(\bar{\theta}, x^*(\alpha_k))} \left[\frac{z}{(\tau_x + \tau_\theta)^{1/2}} + \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta} \right] \left[1 - \Phi \left(\frac{x^*(\alpha_l) - \frac{z}{(\tau_x + \tau_\theta)^{1/2}} - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{\tau_x^{-1/2}} - \alpha_k \right) \right] \phi(z) dz \\ + \int_{L(\bar{\theta}, x^*(\alpha_k))}^{\infty} \left[\frac{z}{(\tau_x + \tau_\theta)^{1/2}} + \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta} \right] \phi(z) dz - T = 0 \end{aligned} \quad (15)$$

where

$$L(\underline{\theta}, x^*(\alpha_k)) = \frac{\theta - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}} \quad \text{and} \quad L(\bar{\theta}, x^*(\alpha_k)) = \frac{\bar{\theta} - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}}$$

We now compute the limit as $\tau_x \rightarrow \infty$ of the LHS of Equation (15). It is straightforward to show that for a sufficiently high τ_x we have $x^*(\alpha_k) \in [T - \delta, \bar{\theta} + \delta]$ for some $\delta > 0$. Thus, it follows that $x^*(\alpha_k, \tau_x)$ has a convergent subsequence. In what follows we work with this subsequence. However, to keep notation simple we slightly abuse notation and talk about limiting behavior $x^*(\alpha_k, \tau_x)$ as $\tau_x \rightarrow \infty$, rather than behavior of $x^*(\alpha_k, \tau_x^n)$ where τ_x^n is an increasing sequence with $\lim_{n \rightarrow \infty} \tau_x^n = \infty$.

Claim 9 *We have*

1. $\lim_{\tau_x \rightarrow \infty} x^*(\alpha_l, \tau_x) = \lim_{\tau_x \rightarrow \infty} x^*(\alpha_k, \tau_x)$ for all $k, l \in \{1, \dots, N\}$.⁶⁸
2. $\lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_l, \tau_x) - x^*(\alpha_k, \tau_x)) \right] \in \mathbb{R}$

Proof. (Part 1) We establish this by contradiction. Wlog suppose that $l < k$. It is straightforward to see that for a sufficiently high τ_x we have $x^*(\alpha_l, \tau_x) > x^*(\alpha_k, \tau_x)$ so that a player who has a higher bias uses a lower threshold. Therefore, $\lim_{\tau_x \rightarrow \infty} x^*(\alpha_l, \tau_x) \geq \lim_{\tau_x \rightarrow \infty} x^*(\alpha_k, \tau_x)$.

Now suppose that $\lim_{\tau_x \rightarrow \infty} x^*(\alpha_l, \tau_x) > \lim_{\tau_x \rightarrow \infty} x^*(\alpha_k, \tau_x)$. It follows that for some $m \in \{1, \dots, N\}$ we have

$$\lim_{\tau_x \rightarrow \infty} x^*(\alpha_m, \tau_x) - \lim_{\tau_x \rightarrow \infty} x^*(\alpha_k, \tau_x) \leq 0 \quad \text{then} \quad \lim_{\tau_x \rightarrow \infty} x^*(\alpha_m, \tau_x) - \lim_{\tau_x \rightarrow \infty} x^*(\alpha_l, \tau_x) < 0$$

⁶⁸Intuitively, if $x^*(\alpha_l) > x^*(\alpha_k)$ then a player with bias α_l when he receives a higher signal not only is sure that a player with bias α_k or lower takes a risky action but also expects to earn $x^*(\alpha_l)$ when the risky action is successful. On the other hand, a player with bias $x^*(\alpha_k)$ can *at most* expect all players with bias α_k or lower take a risky action and expects to receive $x^*(\alpha_k)$ when the risky action is successful. In the limit, this implies that either a player with bias α_l strictly prefer taking the risky action to taking the safe action or a player with bias α_k strictly prefer taking safe action rather than the risky action. Either statement contradicts the definition of threshold signals.

At this point we have to consider two cases: (1) $\lim_{\tau_x \rightarrow \infty} x^*(\alpha_l, \tau_x) < \bar{\theta}$ and (2) $\lim_{\tau_x \rightarrow \infty} x^*(\alpha_l, \tau_x) = \bar{\theta}$.⁶⁹ Since the argument in both cases is analogous we consider only the case when $\lim_{\tau_x \rightarrow \infty} x^*(\alpha_l, \tau_x) < \bar{\theta}$.

If $\lim_{\tau_x \rightarrow \infty} x^*(\alpha_l, \tau_x) < \bar{\theta}$ then

$$\lim_{\tau_x \rightarrow \infty} L(\bar{\theta}, x^*(\alpha_l, \tau_x)) = \infty = \lim_{\tau_x \rightarrow \infty} L(\bar{\theta}, x^*(\alpha_k, \tau_x))$$

Therefore, the indifference condition of a player with bias α_k converges to

$$T = \sum_{m=1}^N g(\alpha_m) \int_{-\infty}^{\infty} x^*(\alpha_k) \left[1 - \Phi \left(-z - \alpha_k + \lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_m, \tau_x) - x^*(\alpha_k, \tau_x)) \right] \right) \right] \phi(z) dz$$

Now, let $n \leq k$ denote the lowest bias such that $\lim_{\tau_x \rightarrow \infty} [x^*(\alpha_k, \tau_x) - x^*(\alpha_n, \tau_x)] \leq 0$, that is for all $\alpha_m < \alpha_n$ we have

$$\lim_{\tau_x \rightarrow \infty} [x^*(\alpha_k, \tau_x) - x^*(\alpha_n, \tau_x)] > 0$$

Then,

$$\begin{aligned} T &= \sum_{m=1}^N g(\alpha_m) \int_{-\infty}^{\infty} x^*(\alpha_k) \left[1 - \Phi \left(-z - \alpha_k + \lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_m, \tau_x) - x^*(\alpha_k, \tau_x)) \right] \right) \right] \phi(z) dz \\ &= \sum_{m=n}^N g(\alpha_m) \int_{-\infty}^{\infty} x^*(\alpha_k) \left[1 - \Phi \left(-z - \alpha_k + \lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_m, \tau_x) - x^*(\alpha_k, \tau_x)) \right] \right) \right] \phi(z) dz \\ &\leq \sum_{m=n}^N g(\alpha_m) \int_{-\infty}^{\infty} x^*(\alpha_k) \phi(z) dz \\ &< \sum_{m=n}^N g(\alpha_m) \int_{-\infty}^{\infty} x^*(\alpha_l) \phi(z) dz \\ &= \sum_{m=n}^N g(\alpha_m) \int_{-\infty}^{\infty} x^*(\alpha_l) \left[1 - \Phi \left(-z - \alpha_l + \lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_m, \tau_x) - x^*(\alpha_k, \tau_x)) \right] \right) \right] \phi(z) dz \\ &\leq \sum_{m=1}^N g(\alpha_m) \int_{-\infty}^{\infty} x^*(\alpha_l) \left[1 - \Phi \left(-z - \alpha_l + \lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_m, \tau_x) - x^*(\alpha_k, \tau_x)) \right] \right) \right] \phi(z) dz = T \end{aligned}$$

where the strict inequality follows from the assumption that $x^*(\alpha_l) > x^*(\alpha_k)$, and the fifth line

⁶⁹In the limit as $\tau_x \rightarrow \infty$ a player will never set a threshold $x^*(\alpha_k)$ strictly above $\bar{\theta}$ since in this case there exists $\varepsilon > 0$ such that $x^*(\alpha_k) - \varepsilon > \bar{\theta}$, and using this lower threshold leads to a strictly higher payoff for all $\theta \in [x^*(\alpha_k) - \varepsilon, x^*(\alpha_k)]$.

follows from observation that

$$\lim_{\tau_x \rightarrow \infty} x^*(\alpha_m, \tau_x) - \lim_{\tau_x \rightarrow \infty} x^*(\alpha_k, \tau_x) \leq 0 \text{ then } \lim_{\tau_x \rightarrow \infty} x^*(\alpha_m, \tau_x) - \lim_{\tau_x \rightarrow \infty} x^*(\alpha_l, \tau_x) < 0$$

Thus, we arrived at a contradiction.

(Part 2) As before we need to differentiate between the case where (1) $\lim_{\tau_x \rightarrow \infty} x^*(\alpha_l, \tau_x) < \bar{\theta}$ and (2) $\lim_{\tau_x \rightarrow \infty} x^*(\alpha_l, \tau_x) = \bar{\theta}$. Since the arguments in both cases are similar, we again only consider here the first case.

We know that $\lim_{\tau_x \rightarrow \infty} x^*(\alpha_l, \tau_x) = \lim_{\tau_x \rightarrow \infty} x^*(\alpha_l, \tau_x)$ for all $l, k \in \{1, \dots, N\}$. Denote this limit by x^* . Now, suppose that there exist k and l such that

$$\lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_k, \tau_x) - x^*(\alpha_l, \tau_x)) \right] \in \{-\infty, +\infty\}$$

Wlog assume that $l < k$. Then we know that for a sufficiently high τ_x we have $x^*(\alpha_k, \tau_x) < x^*(\alpha_l, \tau_x)$. Therefore

$$\lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_k, \tau_x) - x^*(\alpha_l, \tau_x)) \right] < \infty$$

So suppose that $\lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_k, \tau_x) - x^*(\alpha_l, \tau_x)) \right] = -\infty$. Since $x^*(\alpha_k, \tau_x) < x^*(\alpha_l, \tau_x)$ for all $m \neq l, k$ we have

$$\lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_k, \tau_x) - x^*(\alpha_m, \tau_x)) \right] \leq \lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_l, \tau_x) - x^*(\alpha_m, \tau_x)) \right]$$

and

$$\lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_k, \tau_x) - x^*(\alpha_k, \tau_x)) \right] = 0 = \lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_l, \tau_x) - x^*(\alpha_k, \tau_x)) \right]$$

But then, in the limit we have

$$\begin{aligned} T &= \sum_{m=1}^N g(\alpha_m) \int_{-\infty}^{\infty} x^* \left[1 - \Phi \left(-z - \alpha_k - \lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_m) - x^*(\alpha_k)) \right] \right) \right] \phi(z) dz - T \\ &< \sum_{m=1}^N g(\alpha_m) \int_{-\infty}^{\infty} x^* \left[1 - \Phi \left(-z - \alpha_k - \lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_m) - x^*(\alpha_l)) \right] \right) \right] \phi(z) dz - T = T \end{aligned}$$

which is a contradiction. Thus, $\lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_k, \tau_x) - x^*(\alpha_l, \tau_x)) \right] > -\infty$. ■

With the above result we can now compute the limit of $\{x^*(\alpha_k)\}_{k=1}^N$ as $\tau_x \rightarrow 0$.

Lemma 10 *Let x^* be the limit of $\lim_{\tau_x \rightarrow \infty} x^*(\alpha_k, \tau_x) = x^*$. Then,*

$$\begin{aligned} x^* &= \frac{1}{\sum_{m=1}^N g(\alpha_m) \Phi\left(\frac{\alpha_m}{\sqrt{2}}\right)} \quad \text{if} \quad \frac{1}{\sum_{m=1}^N g(\alpha_m) \Phi\left(\frac{\alpha_m}{\sqrt{2}}\right)} \leq \bar{\theta} \\ x^* &= \bar{\theta} \quad \text{otherwise} \end{aligned}$$

Proof. Suppose that

$$\frac{1}{\sum_{m=1}^N g(\alpha_m) \Phi\left(\frac{\alpha_m}{\sqrt{2}}\right)} \leq \bar{\theta}$$

Define

$$\kappa(k, 1) \equiv \lim_{\tau_x \rightarrow \infty} \tau_x^{1/2} [x^*(\alpha_k) - x^*(\alpha_1)] \in \mathbb{R}$$

where the observation that $\kappa(k, 1) \in \mathbb{R}$ follows from Claim 9. As $\tau_x \rightarrow \infty$ the indifference condition for a player with bias α_m , $m \in \{1, \dots, N\}$ converges to

$$x^* \sum_{m=1}^N g(\alpha_m) \int_{-\infty}^{\infty} [1 - \Phi(-z - \alpha_k - \kappa(m, 1) - \kappa(1, k))] \phi(z) dz = 0$$

Note that $\kappa(1, k) = -\kappa(k, 1)$, and therefore the indifference condition is given by

$$x^* \sum_{m=1}^N g(\alpha_m) \int_{-\infty}^{\infty} [1 - \Phi(-z - \alpha_k - \kappa(m, 1) + \kappa(k, 1))] \phi(z) dz - T = 0$$

or

$$x^* \sum_{m=1}^N g(\alpha_m) \Phi\left(\frac{\alpha_k + \kappa(m, 1) + \kappa(k, 1)}{\sqrt{2}}\right) - T = 0$$

Thus, we arrive at the following system of $N - 1$ unknowns (the unknown being $\{\kappa(m, 1)\}_{m=1}^N$):

$$\begin{aligned} x^* \sum_{m=1}^N g(\alpha_m) \Phi\left(\frac{\alpha_1 + \kappa(m, 1) + \kappa(1, 1)}{\sqrt{2}}\right) - T &= 0 \\ &\vdots \\ x^* \sum_{m=1}^N g(\alpha_m) \Phi\left(\frac{\alpha_N + \kappa(m, 1) + \kappa(N, 1)}{\sqrt{2}}\right) - T &= 0 \end{aligned}$$

where the first equation corresponds to the indifference condition of a player with bias α_1 in the limit, the second equation corresponds to the indifference condition of a player with bias α_2 and so

on. By inspection, one can verify that the solution to this equation is

$$\kappa(m, 1) = \alpha_m - \alpha_1, \text{ for all } m \in \{1, \dots, N\}$$

Moreover, this is a unique solution. To see this, note that equations 2 to N above are a set of $N - 1$ non-linear equations in $N - 1$ unknowns.⁷⁰ It is straightforward to show that the Jacobian of this system is diagonally dominant which, by Lemma 8 implies that this system has a unique solution. It follows that

$$x^* = \frac{T}{\sum_{m=1}^N g(\alpha_m) \Phi\left(\frac{\alpha_m}{\sqrt{2}}\right)}$$

If $T > \bar{\theta} \sum_{m=1}^N g(\alpha_m) \Phi(\alpha_m/\sqrt{2})$ then x^* converges to $\bar{\theta}$. We prove this by contradiction. Suppose that $x^*(\alpha_k)$'s do not converge to $\bar{\theta}$. Since we know that $\lim_{\tau_x \rightarrow \infty} x^*(\alpha_k) \in [T, \bar{\theta}]$ it follows that in this case we must have $\lim_{\tau_x \rightarrow \infty} x^*(\alpha_k) < \bar{\theta}$ and so

$$\bar{L} = \lim_{\tau_x} \frac{\bar{\theta} - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}} = \infty$$

We perform the substitution $z = \left[\theta - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}\right] / (\tau_x + \tau_\theta)^{-1/2}$ in the indifference equation of a player with bias α_k , $k \in \{1, \dots, N\}$ and take the limit as $\tau_x \rightarrow \infty$ to obtain

$$\sum_{m=1}^N g(\alpha_m) \int_{-\infty}^{\infty} x^* [1 - \Phi(-z - \alpha_k - \kappa(m, 1) - \kappa(1, k))] \phi(z) dz = 0,$$

where

$$\kappa(k, l) = \lim_{\tau_x \rightarrow \infty} \tau_x^{1/2} [x^*(\alpha_k) - x^*(\alpha_l)] \in \mathbb{R}$$

and so $\kappa(k, l) = \kappa(k, 1) - \kappa(1, l) = \kappa(k, 1) + \kappa(l, 1)$. Evaluating the resulting integrals, we arrive at the system of equations:

$$\begin{aligned} x^* \sum_{m=1}^N g(\alpha_m) \Phi\left(\frac{\alpha_1 + \kappa(m, 1) + \kappa(1, 1)}{\sqrt{2}}\right) - T &= 0 \\ &\vdots \\ x^* \sum_{m=1}^N g(\alpha_m) \Phi\left(\frac{\alpha_N + \kappa(m, 1) + \kappa(N, 1)}{\sqrt{2}}\right) - T &= 0 \end{aligned}$$

⁷⁰There are only $N - 1$ independent constants $\kappa(k, l)$ since $\kappa(k, k) = 0$, $\kappa(k, l) = \kappa(l, k)$, and $\kappa(k, l) = \kappa(k, m) + \kappa(m, l)$.

But we showed above that this system has a unique solution

$$x^* = \frac{T}{\sum_{m=1}^N g(\alpha_m) \Phi\left(\frac{\alpha_m}{\sqrt{2}}\right)} > \bar{\theta}$$

by assumption. This is a contradiction since x^* was supposed to be smaller than $\bar{\theta}$. It follows that $x^* = \bar{\theta}$. ■

The argument presented above establishes that any convergent sequence $x^*(\alpha_k, \tau_x^n) \rightarrow x^*$, where $x^* = T / [\sum_{m=1}^N g(\alpha_m) \Phi(\alpha_m/\sqrt{2})]$ if $T / [\sum_{m=1}^N g(\alpha_m) \Phi(\alpha_m/\sqrt{2})] < \bar{\theta}$ and $x^* = \bar{\theta}$ otherwise. But then from Lemma it follows that $x^*(\alpha_k, \tau_x) \rightarrow x^*$. Thus, we established Part (2) of Proposition 3

D.2 Uniqueness of equilibrium away from the limit

In this section we prove part 1 of Proposition 3, that is we show that there exists $\bar{\tau}_x$ such that for all $\tau_x > \bar{\tau}_x$ the model with biased beliefs has unique equilibrium.

Lemma 11 *There exists $\bar{\tau}_x$ such that if $\tau_x > \bar{\tau}_x$ then the model has unique equilibrium in monotone strategies.*

Proof. To establish this result it is enough to show that for a sufficiently high τ_x

$$\frac{\partial V_k}{\partial x^*(\alpha_k)} - \sum_{l \neq k} \left| \frac{\partial V_l}{\partial x^*(\alpha_k)} \right| > 0 \text{ for all } k \in \{1, \dots, N\},$$

which implies that the Jacobian of the mapping V is diagonally dominant. Using the expressions for the derivatives reported above we have

$$\begin{aligned}
& \frac{\partial V_k}{\partial x^*(\alpha_k)} - \sum_{l \neq k} \frac{\partial V_l}{\partial x^*(\alpha_k)} \\
= & \sum_{l=1}^N g(\alpha_l) \frac{\tau_x}{\tau_x + \tau_\theta} \int_{\underline{\theta}}^{\bar{\theta}} \left[1 - \Phi \left(\frac{x^*(\alpha_l) - \theta}{\tau_x^{-1/2}} - \alpha_k \right) \right] (\tau_x + \tau_\theta)^{1/2} \phi \left(\frac{\theta - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}} \right) d\theta \\
& + \sum_{l=1}^N g(\alpha_l) \frac{\tau_x}{\tau_x + \tau_\theta} \int_{\underline{\theta}}^{\bar{\theta}} \theta \tau_x^{1/2} \phi \left(\frac{x^*(\alpha_l) - \theta}{\tau_x^{-1/2}} - \alpha_k \right) (\tau_x + \tau_\theta)^{1/2} \phi \left(\frac{\theta - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}} \right) d\theta \\
& - g(\alpha_k) \frac{\tau_\theta}{\tau_x + \tau_\theta} \int_{\underline{\theta}}^{\bar{\theta}} \theta \tau_x^{1/2} \phi \left(\frac{x^*(\alpha_l) - \theta}{\tau_x^{-1/2}} - \alpha_k \right) (\tau_x + \tau_\theta)^{1/2} \phi \left(\frac{\theta - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}} \right) d\theta \\
& - \sum_{l \neq k} g(\alpha_l) \int_{\underline{\theta}}^{\bar{\theta}} \theta \tau_x^{1/2} \phi \left(\frac{x^*(\alpha_l) - \theta}{\tau_x^{-1/2}} - \alpha_k \right) (\tau_x + \tau_\theta)^{1/2} \phi \left(\frac{\theta - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}} \right) d\theta \\
& + \text{Positive Terms}
\end{aligned}$$

where $\sum_{l \neq k} \left| \frac{\partial V_l}{\partial x^*(\alpha_k)} \right|$ in the above expression is captured by

$$\sum_{l \neq k} g(\alpha_l) \int_{\underline{\theta}}^{\bar{\theta}} \theta \tau_x^{1/2} \phi \left(\frac{x^*(\alpha_l) - \theta}{\tau_x^{-1/2}} - \alpha_k \right) (\tau_x + \tau_\theta)^{1/2} \phi \left(\frac{\theta - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}} \right) d\theta$$

We focus on the first four terms in the above derivative and notice that they simplify to

$$\begin{aligned}
& \sum_{l=1}^N g(\alpha_l) \frac{\tau_x}{\tau_x + \tau_\theta} \int_{\underline{\theta}}^{\bar{\theta}} \left[1 - \Phi \left(\frac{x^*(\alpha_k) - \theta}{\tau_x^{-1/2}} - \alpha_k \right) \right] (\tau_x + \tau_\theta)^{1/2} \phi \left(\frac{\theta - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}} \right) d\theta \\
& - \sum_{l=1}^N g(\alpha_l) \frac{\tau_\theta}{\tau_x + \tau_\theta} \int_{\underline{\theta}}^{\bar{\theta}} \theta \tau_x^{1/2} \phi \left(\frac{x^*(\alpha_l) - \theta}{\tau_x^{-1/2}} - \alpha_k \right) (\tau_x + \tau_\theta)^{1/2} \phi \left(\frac{\theta - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}} \right) d\theta
\end{aligned}$$

To show that the above term is strictly positive for a high τ_x , in each of the above integrals we make the substitution

$$z = \frac{\theta - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}}$$

Then, the above integrals become

$$\begin{aligned}
& \sum_{l=1}^N g(\alpha_l) \frac{\tau_x}{\tau_x + \tau_\theta} \int_{L(\underline{\theta}, \alpha_k)}^{L(\bar{\theta}, \alpha_k)} \left[1 - \Phi \left(\frac{x^*(\alpha_k) - \frac{z}{(\tau_x + \tau_\theta)^{1/2}} - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{\tau_x^{-1/2}} - \alpha_k \right) \right] \phi(z) dz \\
& - \sum_{l=1}^N g(\alpha_l) \frac{\tau_\theta}{\tau_x + \tau_\theta} \int_{L(\underline{\theta}, \alpha_k)}^{L(\bar{\theta}, \alpha_k)} \left\{ \left[\frac{z}{(\tau_x + \tau_\theta)^{1/2}} + \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta} \right] \times \right. \\
& \left. \tau_x^{1/2} \phi \left(\frac{x^*(\alpha_k) - \frac{z}{(\tau_x + \tau_\theta)^{1/2}} - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{\tau_x^{-1/2}} - \alpha_k \right) \phi(z) \right\} dz
\end{aligned}$$

where

$$L(\bar{\theta}, \alpha_k) = \frac{\bar{\theta} - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}} \quad \text{and} \quad L(\underline{\theta}, \alpha_k) = \frac{\underline{\theta} - \frac{\tau_x x^*(\alpha_k) + \tau_\theta \mu_\theta}{\tau_x + \tau_\theta}}{(\tau_x + \tau_\theta)^{-1/2}}$$

Since $\lim x^*(\alpha_k) \leq \bar{\theta}$ we have $L(\underline{\theta}, \alpha_k) \rightarrow -\infty$ as $\tau_x \rightarrow \infty$ and that $\lim_{\tau_x \rightarrow \infty} L(\bar{\theta}, \alpha_k) > -\infty$.

Therefore,

$$\begin{aligned}
& \lim_{\tau_x \rightarrow \infty} \left[\frac{\partial V_k}{\partial x^*(\alpha_k)} - \sum_{l \neq k} \frac{\partial V_l}{\partial x^*(\alpha_k)} \right] \\
& \geq \sum_{l=1}^N g(\alpha_l) \int_{-\infty}^{\bar{L}} \left[1 - \Phi \left(-z - \alpha_k + \lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_l) - x^*(\alpha_k)) \right] \right) \right] \phi(z) dz \\
& - \left(\lim_{\tau_x \rightarrow \infty} \frac{\tau_\theta}{\tau_x + \tau_\theta} \right) \sum_{l=1}^N g(\alpha_l) \int_{-\infty}^{\bar{L}} x^* \tau_x^{1/2} \phi \left(-z - \alpha_k + \lim_{\tau_x \rightarrow \infty} \left[\tau_x^{1/2} (x^*(\alpha_l) - x^*(\alpha_k)) \right] \right) \phi(z) \\
& > 0
\end{aligned}$$

It follows that there exists a large enough τ_x , call it $\bar{\tau}_{x,k}$, such that if $\tau_x > \bar{\tau}_{x,k}$ then the k -th row of the Jacobian of mapping \mathbf{V} is dominated by its diagonal entry. Define

$$\bar{\tau}_x = \max_{k \in \{1, \dots, N\}} \bar{\tau}_{x,k}$$

Then for all $\tau_x \geq \bar{\tau}_x$ the Jacobian of mapping \mathbf{V} is diagonally dominant, which, by Lemma 8, establishes the claim. ■

References

- [1] Avoyan, A (2018). Communication in Global Games: Theory and Experiment. Working paper
- [2] Alaoui, L., and Penta, A. (2015). Endogenous depth of reasoning. *Review of Economic Studies*, 83(4), 1297-1333.
- [3] Angeletos, G. M., and Lian, C. (2016). Incomplete information in macroeconomics: Accommodating frictions in coordination. *Handbook of Macroeconomics*, 2, 1065-1240.
- [4] Angeletos, G. M. and La'O, J. (2013). Sentiments. *Econometrica*, 81(2), 739-780.
- [5] Angeletos, G. M., Hellwig, C., & Pavan, A. (2006). Signaling in a global game: Coordination and policy traps. *Journal of Political Economy*, 114(3), 452-484.
- [6] Ash, R. (2000). Probability and measure theory. Academic Press.
- [7] Bannier, C. E., and Heinemann, F. (2005). Optimal transparency and risk-taking to avoid currency crises. *Journal of Institutional and Theoretical Economics*, 161(3), 374-391.
- [8] Baeriswyl, R., and Cornand, C. (2014). Reducing overreaction to central banks' disclosures: Theory and experiment. *Journal of the European Economic Association*, 12(4), 1087-1126.
- [9] Benhabib, J. , P. Wang, and Y. Wen (2015). Sentiments and aggregate demand fluctuations. *Econometrica*, 83(2), 549-585.
- [10] Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3), 623-685.
- [11] Brandts, J., and Charness, G. (2011). The strategy versus the direct-response method: a first survey of experimental comparisons. *Experimental Economics*, 14(3), 375-398.
- [12] Brandenburger, A. (1996). Strategic and Structural Uncertainty in Games. In R. Zeckhauser, R. Keeney and J. Sibenius (eds.) *Wise Choices: Games, Decisions, and Negotiations* (Boston: Harvard Business School Press), 221–232.
- [13] Cabrales, A., Garcia-Fontes, W., and Motta, M. (2000). Risk dominance selects the leader: An experimental analysis. *International Journal of Industrial Organization*, 18(1), 137-162.
- [14] Cabrales, A., Nagel, R., and Armenter, R. (2007). Equilibrium selection through incomplete information in coordination games: an experimental study. *Experimental Economics*, 10(3), 221-234.

- [15] Camerer, C., and Lovo, D. (1999). Overconfidence and excess entry: An experimental approach. *American Economic Review*, 89(1), 306-318.
- [16] Carlsson, H. and Van Damme, E., 1993. Global games and equilibrium selection. *Econometrica* 61(5), 989-1018.
- [17] Charness, G., Feri, F., Meléndez-Jiménez, M. A., and Sutter, M. (2014). Experimental games on networks: Underpinnings of behavior and equilibrium selection. *Econometrica*, 82(5), 1615-1670.
- [18] Colombo, L., Femminis, G., and Pavan, A. (2014). Information acquisition and welfare. *Review of Economic Studies*, 81(4), 1438-1483.
- [19] Cooper, R. W., DeJong, D. V., Forsythe, R., and Ross, T. W. (1990). Selection criteria in coordination games: Some experimental results. *American Economic Review*, 80(1), 218-233.
- [20] Cooper, R., DeJong, D. V., Forsythe, R., and Ross, T. W. (1992). Communication in coordination games. *The Quarterly Journal of Economics*, 107(2), 739-771.
- [21] Cornand, C., and Heinemann, F. (2014). Measuring agents' reaction to private and public information in games with strategic complementarities. *Experimental Economics*, 17(1), 61-77.
- [22] Costa-Gomes, M. A., and Crawford, V. P. (2006). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review*, 96(5), 1737-1768.
- [23] Croson, R. T. (2000). Thinking like a game theorist: factors affecting the frequency of equilibrium play. *Journal of Economic Behavior & Organization*, 41(3), 299-314.
- [24] Darai, D., Kogan, S., Kwasnica, A., and Weber, R. (2017) "Aggregate Sentiment and Investment: An Experimental Study" mimeo.
- [25] Duffy, J., and Ochs, J. (2012). Equilibrium selection in static and dynamic entry games. *Games and Economic Behavior*, 76(1), 97-116.
- [26] Diamond, D. W. and Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91 (3): 401-419.
- [27] Edmond, C. (2013). Information manipulation, coordination, and regime change. *Review of Economic Studies*, 80(4), 1422-1458.
- [28] Ely, J., and Peski, M. (2006). Hierarchies of belief and interim rationalizability. *Theoretical Economics* 1 (2006): 19-65.

- [29] Eyster, E., and Rabin, M. (2005). Cursed equilibrium. *Econometrica*, 73(5), 1623-1672.
- [30] Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171-178.
- [31] Frankel, D. M., Morris, S., and Pauzner, A. (2003). Equilibrium selection in global games with strategic complementarities. *Journal of Economic Theory*, 108(1), 1-44.
- [32] Gächter, S., and Renner, E. (2010). The effects of (incentivized) belief elicitation in public goods experiments. *Experimental Economics*, 13(3), 364-377.
- [33] Goldstein, I., and Pauzner, A. (2005). Demand–deposit contracts and the probability of bank runs. *Journal of Finance*, 60(3), 1293-1327.
- [34] García, D., Sangiorgi, F., and Urošević, B. (2007). Overconfidence and market efficiency with heterogeneous agents. *Economic Theory*, 30(2), 313-336.
- [35] Goeree, J., Holt, C., and Palfrey, T. (2016). *Quantal Response Equilibrium: A Stochastic Theory of Games*. Princeton University Press.
- [36] Harsanyi, J. and R. Selten (1988). “A General Theory of Equilibrium Selection in Games,” The MIT Press.
- [37] Heinemann, F. and G. Illing (2002). Speculative attacks: unique equilibrium and transparency. *Journal of International Economics*, 58(2): 429-450.
- [38] Heinemann, F., Nagel, R., and Ockenfels, P. (2004). The theory of global games on test: experimental analysis of coordination games with public and private information. *Econometrica*, 72(5), 1583-1599.
- [39] Heinemann, F., Nagel, R., and Ockenfels, P. (2009). Measuring strategic uncertainty in coordination games. *Review of Economic Studies*, 76(1), 181-221.
- [40] Hellwig, C. (2002). Public information, private information, and the multiplicity of equilibria in coordination games. *Journal of Economic Theory*, 107(2), 191-222.
- [41] Hellwig, C., Mukherji, A., and Tsyvinski, A. (2006). Self-fulfilling currency crises: The role of interest rates. *American Economic Review*, 1769-1787.
- [42] Jehiel, P., 2005. Analogy-based expectation equilibrium. *Journal of Economic Theory*, 123(2): 81-104.

- [43] Jehiel, P., and Koessler, F. (2008). Revisiting games of incomplete information with analogy-based expectations. *Games and Economic Behavior*, 62(2), 533-557.
- [44] Iachan, F. and P. Nenov (2015). Information quality and crises in regime-change games. *Journal of Economic Theory*, 158: 739-768.
- [45] Izmalkov, S. and Yildiz, M. (2010). Investor sentiments. *American Economic Journal: Microeconomics*, 2(1): 21-38.
- [46] Hellwig, C., and Veldkamp, L. (2009). Knowing what others know: Coordination motives in information acquisition. *Review of Economic Studies*, 76(1), 223-251.
- [47] Kagel, J. H., and Roth, A. E. (Eds.). (2016). The Handbook of Experimental Economics, Volume 2: The Handbook of Experimental Economics. Princeton University Press.
- [48] Kneeland, T. (2016). Coordination under limited depth of reasoning. *Games and Economic Behavior*, 96, 49-64.
- [49] McKelvey, R. D., and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1), 6-38.
- [50] Moore, D. and Cain, D. (2007). Overconfidence and underconfidence: When and why people underestimate (and overestimate) the competition. *Organizational Behavior and Human Decision Processes*, 103: 197-213.
- [51] Morris, S., and Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 88(3), 587-597.
- [52] Morris, S., and Shin, H. S. (2002). Measuring strategic uncertainty. Mimeo.
- [53] Morris, S., and Shin, H. S. (2001). Global games: Theory and applications. In *Advances in Economics and Econometrics: Theory and applications*, Eight World Congress, Cambridge University Press.
- [54] Morris, S., Shin, H. S., and Yildiz, M. (2016). Common belief foundations of global games. *Journal of Economic Theory*, 163, 826-848.
- [55] Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5), 1313-1326.
- [56] Pavan, A. (2016). Attention, Coordination, and Bounded Recall. Working paper
- [57] Persico, N. (2000). Information acquisition in auctions. *Econometrica*, 68(1), 135-148.

- [58] Rutström, E. E., and Wilcox, N. T. (2009). Stated beliefs versus inferred beliefs: A methodological inquiry and experimental test. *Games and Economic Behavior*, 67(2), 616-632.
- [59] Stahl, D. O., and Wilson, P. W. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1), 218-254.
- [60] Schotter, A., and Trevino, I. (2015). Belief elicitation in the laboratory. *Annual Review of Economics*, 6(1), 103-128.
- [61] Straub, P. G. (1995). Risk dominance and coordination failures in static games. *The Quarterly Review of Economics and Finance*, 35(4), 339-363.
- [62] Strzalecki, T. (2014). Depth of reasoning and higher order beliefs. *Journal of Economic Behavior & Organization*, 108, 108-122.
- [63] Szkup, M. (2016). Preventing self-fulfilling debt crises. UBC working paper
- [64] Szkup, M., and Trevino, I. (2015). Information acquisition in global games of regime change. *Journal of Economic Theory*, 160, 387-428.
- [65] Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review*, 80(1), 234-248.
- [66] Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1991). Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. *The Quarterly Journal of Economics*, 106(3), 885-910.
- [67] Van Zandt, T., and Vives, X. (2007). Monotone equilibria in Bayesian games of strategic complementarities. *Journal of Economic Theory*, 134(1), 339-360.
- [68] Van Zandt, T. (2010). Interim Bayesian Nash equilibrium on universal type spaces for supermodular games. *Journal of Economic Theory*, 145(1), 249-263.
- [69] Vives, X. (2010) Information and learning in markets: the impact of market microstructure. Princeton University Press.
- [70] Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, 39(5), 806.
- [71] Yang, M. (2015). Coordination with flexible information acquisition. *Journal of Economic Theory*, 158, 721-738.