

Expectation Conformity in Strategic Cognition*

Alessandro Pavan[†] and Jean Tirole[‡]

PRELIMINARY AND INCOMPLETE

Abstract

The paper studies “cognitive games,” that is, games in which the players can influence their understanding of a strategic situation before playing the primitive (normal- or extensive-form) game. The analysis covers both the case of self-directed cognition (as when a player controls her own information structure) and the case of manipulative cognition (as when a player influences her opponents’ understanding of the game). We introduce the concept of *expectation conformity* and show how the latter, together with its decomposition into *unilateral expectation conformity* and *increasing differences*, sheds light on the choice of the cognitive structures (both on and off the equilibrium path) and on the sensitivity of the cognitive postures to the type of strategic interaction (e.g., complements vs substitutes). We show that constant-sum games never give rise to self-fulfilling cognition. By contrast, the latter emerges in many non-constant-sum games, both when cognition is self-directed and takes the form of “sparsity,” noisy information acquisition, or “espionage,” (i.e., learning about others’ beliefs), as well as when cognition is manipulative and takes the form of framing, signal jamming, noisy disclosures, and counter-intelligence. Finally, we discuss the role that expectation conformity plays in games with boundedly-rational players such as those considered in the level-k literature.

Keywords: cognition, expectation conformity, cognitive traps, sparsity, framing, memory management, endogenous depth of reasoning.

JEL numbers: C72; C78; D82; D83; D86.

*Research support from European Research Council advanced grant (European Community’s Seventh Framework Programme (FP7/2007-2013) Grant Agreement no. 249429 and European Union’s Horizon 2020 research and innovation programme, Grant Agreement no. 669217) is gratefully acknowledged. We are grateful to Matteo Camboni for excellent research assistance.

[†]Department of Economics, Northwestern University, and CEPR. Email: alepavan@northwestern.edu.

[‡]Toulouse School of Economics (TSE) and Institute for Advanced Study in Toulouse (IAST), University of Toulouse Capitole. Email: jean.tirole@tse-fr.eu.

1 Introduction

Cognition is costly. Thinking about possible contingencies, memorizing information, forming beliefs about other players' understanding of a game as a way to predict their behavior, but also brainstorming with others and obtaining the financial, engineering, and legal expertise necessary to play a game requires non-trivial investments, whose cost depends on the urgency to act, the cognitive load, and the context.

In this paper, we consider strategic situations in which players' cognition (that is, their understanding of the game) is endogenous. Depending on the context under examination, cognition may take the form of information acquisition about a payoff-relevant state, learning about other players' information, manipulating other players' beliefs (for example, through signal jamming, or by selecting appropriate frames), memory management, but also iterating over best responses to identify the payoff-maximizing actions.

We are particularly interested in how a player's cognitive posture depends on (a) her beliefs about other players' expectations about her own cognitive posture, and (b) her beliefs about other players' actual cognition. We are also interested in understanding how the cognitive choices depend on the type of strategic situation, in particular on whether actions in the primitive game are strategic complements or substitutes. The analysis permits us to determine whether cognitive choices are themselves "strategic complements" or "strategic substitutes" and whether or not the game admits multiple equilibria with different cognitive investments, without having to solve for the actual equilibrium cognitive choices, which may be daunting in many situations of interest.

To fix ideas, suppose that the players' cognitive choices are indexed by $(\rho_i)_{i \in I}$, where I is the set of players. Such cognitive choices may have different interpretations depending on the strategic and cognitive situation under consideration. We come back to these specific interpretations in due course. For the time being, interpret such cognitive activities flexibly as investments that influence the players' understanding of the primitive downstream normal- or extensive-form game.

We say that *expectation conformity* (EC) holds for the cognitive profiles $\rho = (\rho_i)_{i \in I}$ and $\hat{\rho} = (\hat{\rho}_i)_{i \in I}$ if each player i 's incentives to move from ρ_i to $\hat{\rho}_i$ are stronger when other players expect him to choose cognition $\hat{\rho}_i$ instead of ρ_i , and when other players themselves choose cognition $\hat{\rho}_{-i}$ instead of cognition ρ_{-i} . We show that EC originates in the interaction between two forces. The first one, *unilateral expectation conformity* (UEC), operates through the influence exerted by the expectations of players other than i about player i 's cognition on player i 's cognitive choice, keeping these other players' cognitive choices fixed. For example, when cognition is *self-directed* (that is, it influences a player's own understanding of the game without affecting others') and takes the form of information acquisition about an exogenous payoff state (with the various information structures ordered a' la Blackwell), UEC obtains when player i has stronger incentives to acquire more information when other players expect her to do so. Alternatively, when cognition is *manipulative* (that is, it influences other players' understanding of the game, as when a player engages in noisy information disclosures, signal jamming, framing, and counter-intelligence), UEC obtains when a player's incentives to manipulate other players' information are stronger, the more other players expect player i to manipulate their information, as in the career-concerns literature (see, e.g., Holmstrom, 1999, and Dewatripont et al. 1999), but also in games of framing, and in many other models of signal jamming. Clearly, UEC is equivalent to EC in games where only one player engages in cognition.

When multiple players engage in cognition, EC can also result from player i 's beliefs about her opponents' actual cognitive choices. We label this second force *increasing differences* (ID) because it is related to the corresponding notion in supermodular games. Fixing the other players' expectations of player i 's cognitive posture, ID holds if player i 's gross payoff from moving from cognition ρ_i to cognition $\hat{\rho}_i$ is larger when other players' cognition is $\hat{\rho}_{-i}$ than when it is ρ_{-i} . Consider again the case in which cognition takes the form of self-directed information acquisition and information structures are ordered, with $\hat{\rho} = (\hat{\rho}_i)_{i \in I}$ denoting a collection of information structures such that each $\hat{\rho}_i$ is Blackwell more informative than ρ_i , where $\rho = (\rho_i)_{i \in I}$ is a different cognitive profile. In this case, ID holds if the gross value to player i of choosing the more informative signal structure $\hat{\rho}_i$ is larger when other players themselves choose the more informative structures $\hat{\rho}_{-i}$. Importantly, the notion of ID that we consider extends to settings where cognition takes forms other than information acquisition, and where the cognitive choices need not have a lattice structure.

We show that EC plays an important role for the determinacy of equilibria and for whether cognitive choices are themselves strategic complements or substitutes. In particular, we show that equilibrium uniqueness obtains, no matter the cognitive costs, if EC is never satisfied across any pair of cognitive profiles $\rho = (\rho_i)_{i \in I}$ and $\hat{\rho} = (\hat{\rho}_i)_{i \in I}$. In contrast, EC leads to equilibrium multiplicity for appropriate cognitive costs. When this is the case, cognitive trap obtains when two ordered equilibria co-exist, and players are better-off in the low-cognition equilibrium.

Equipped with the above results, we then specialize the analysis to specific games, as well as to settings in which cognition takes specific forms. We start by considering a class of downstream interactions in which, in each state, the sum of the players' payoffs is constant. We show that, in such games, EC is either violated or holds as an equality, meaning that, if multiple equilibria exist, a player is indifferent over any of her equilibrium cognitive levels.

In contrast, EC emerges in many non-constant-sum games. First, we consider games in which cognition is *self-directed* and takes the form of "sparsity" (see Gabaix, 2014, for an introduction to sparsity in non-strategic settings). Sparsity holds when a decision maker decides to pay attention only to a few dimensions of a relevant payoff state, and then reasons and acts as if any of the non-explored dimensions did not exist. There is a natural progression in reasoning, whereby certain dimensions must be explored before others. In such situations, a player's depth of cognition ρ_i corresponds to the number of dimensions explored.¹ Sparsity is typically intended to capture bounded rationality. We show that its key features are also consistent with a certain representation where all players are fully rational. In a strategic setting, sparsity comes with interesting properties. For example, a player who goes "deeper" in the exploration of the game can perfectly predict the behavior of any player who explores fewer dimensions, whereas a player who explores fewer dimensions than her opponents reasons and acts as if the opponents explored the same dimensions that she explored, despite knowing that this is not the case. We relate EC to whether downstream actions are strategic complements or substitutes and then use EC to identify various features of the equilibrium set. For example, we show that all pure-strategy equilibria are necessarily symmetric in the complements case, whereas, in the substitutes case, there exists at most one symmetric pure-strategy equilibrium, typically co-existing with many asymmetric pure-strategy equilibria. Furthermore, while, in the complements case, total welfare is maximal in the symmetric equilibrium featuring the largest cognition, in the substitute

¹We also discuss instances in which the decision maker chooses the order by which she explores the various dimensions.

case, it is maximal in the asymmetric equilibrium featuring the lowest cognition for the player who is behind in the exploration of the state.

Remaining in the realm of self-directed cognition, we then move to games in which cognition takes the perhaps more familiar form of noisy information acquisition, as when each player chooses the precision of an additive signal about an unknown payoff state. We show that, no matter whether actions in the primitive game are strategic complements or substitutes, UEC always holds in these games. ID, instead, holds for strategic complements but not for substitutes.

Cognition in most of the literature on information acquisition takes the form of players learning about payoff states. We also consider the case where players acquire information about other players' beliefs about payoffs. Examples include situations where players spy on other players, as in certain industrial-espionage games, but also situations where players invest in understanding other players' culture, cognitive traits, and other dimensions of their personality that are responsible for their view of the game. We capture such situations by letting the players choose the precision of a signal they receive about other players' signals. UEC always holds in these games, irrespective of whether downstream actions are complements or substitutes. ID, instead, is reversed relative to the case where players learn about a payoff state: a player's incentives to spy on her opponents are stronger when her opponents themselves are expected to spy more and the downstream actions are substitutes, whereas they are weaker when the opponents are expected to spy more and the downstream actions are complements.

Next, we consider situations in which cognition is *manipulative*, that is, it affects other players' understanding of the game. First, we consider games where the manipulation operates through frames influencing other players' recollection of previously received information. Specifically, we consider a class of information-design games in which the persuader (the Sender) has state-invariant preferences and benefits from inducing the other player (the Receiver) to take a higher action. The Receiver, instead, has state-dependent preferences and wants to take a high action in high states and a low action in low states.² Contrary to what typically assumed in the information-design literature, we assume that the Sender cannot commit to her choice of a manipulative structure. We also allow the Receiver to engage in defensive memory management, that is, in various activities that help her recollect information previously received, thus reducing the effect of the Sender's manipulation on her decision. We show that, as in other signal-jamming environments (most notably, those considered in the career-concerns literature), UEC always holds for the Sender. When the Receiver expects the Sender to engage in more manipulation, she interprets her inability to recollect information favorable to the Sender as informative of the state being unfavorable to the Sender, which in turn makes it particularly profitable to the Sender to engage in manipulation to avoid such situations.³ Whether or not such games also feature ID is in general more convoluted but we are able to identify certain conditions and relate them to the primitives of the game.

We also consider games in which cognition takes the form of players sharing their information with others (as in certain noisy communication games), or in obstructing others from spying on their

²As discussed in the information-design literature, many strategic situations fit such a description (think of a seller trying to persuade a buyer to expand her purchases, or a policymaker trying to persuade his constituency to renew a political endorsement).

³In these games, UEC holds as an equality for the Receiver. The latter is indifferent as to what expectations the Sender holds about her cognitive posture for, in such games, only the Receiver acts in the downstream game.

information (as when a firm invests in counter-intelligence to make it difficult for rival firms to scoop on her research output). Contrary to the case where cognition is self-directed, we show that UEC holds when downstream actions are complements but not when they are substitutes. Importantly, these games feature a negative form of ID: a player’s incentives to share information with others, or to let other players spy on her, are stronger, the less information other players share, or the more they invest in counter-intelligence to prevent others from spying on them. This is true irrespective of whether downstream actions are complements or substitutes.

In all of the various settings discussed above, cognition takes the form of the choice of an information structure (for oneself, when self-directed, or for other players, when manipulative). We also consider situations in which payoffs are commonly known and cognition instead affects a player’s ability to compute iterated best responses. Specifically, we endogenize the depth of reasoning in a model of level- k thinking where the depth of reasoning defines the maximal number of iterations over best responses a player is able to perform. Our analysis is inspired by recent theoretical and experimental work by Alaoui and Penta (2016, 2017, 2018) for the level- k model (see Crawford, Costa-Gomes, and Iriberri, 2013, for a review of the level- k literature). The key point of departure from previous work is in accommodating for the possibility that a player’s cognitive choice (her depth of reasoning) depends on her beliefs about her opponents’ cognitive level (ID) as well as on her beliefs about her opponents’ expectations about her own depth of reasoning (UEC). In such games, each player goes at most one level deeper than the opponent in the computation of the best responses. We consider a downstream game first introduced in Arad and Rubinstein (2012) in which a player is rewarded for minimally undercutting her rivals. In the case of asymmetric cognitive profiles, UEC holds as equality for both the leader (the player going deeper in the computation of best responses) and the follower, meaning that neither player responds to variations in the opponent’s expectations of her own cognition. More interestingly, the leader’s incentives to go deeper decrease with the follower’s depth of reasoning, whereas the opposite is true for the follower. When, instead, the two players are expected to engage in the same cognition, their incentives to go deeper in the exploration of the best responses are always smaller, the larger the opponent’s depth of reasoning — a form of negative ID. When taken together, the above results imply that this game never admits multiple equilibria, irrespective of the cognitive costs.

Summarizing, we view the paper’s contribution as twofold. First, EC, and its decomposition into UEC and ID, provides a powerful tool to examine the players’ incentives to engage in various forms of cognition, which in turn helps connecting the determinacy of equilibria and the cognitive postures to the type of strategic interactions in the primitive game. Second, the paper provides a unifying perspective over the forces that shape the cognitive choices. By bringing apparently distinct phenomena under the same conceptual umbrella, the paper identifies common themes but also isolates the distinctive role played by specific forms of cognition.

Organization. The rest of the paper is organized as follows. We wrap up the Introduction below with a brief discussion of the most pertinent literature. Section 2.1 contains the description of cognitive games in which cognition takes the form of the choice of an information structure. It also contains the definition of EC, UEC and ID, and the results relating these concepts to the determinacy of equilibria. Section 3 contains the results for constant-sum games. Section 4 contains the analysis of games in which cognition is self-directed and takes the form of sparsity, noisy information acquisition about payoff states, and espionage (i.e., noisy information acquisition about other players’ beliefs).

Section 5 contains the results for games in which cognition is manipulative and takes the form of framing, generalized career-concerns, noisy information sharing, and counter-intelligence. Section 6 contains the results for cognitive games in which payoffs are known and cognition determines the ability to compute best responses, as in the level-k model. Section 7 concludes. All omitted proofs are in the Appendix at the end of the document.

Related Literature: The paper is related to a few strands of the literature. The first one is the literature on information acquisition in strategic settings. This literature traces back at least to Stigler (1961)'s analysis of search models. More recently, information acquisition has been studied in global games (e.g., Szkup and Trevino, 2015, Yang, 2015, and Morris and Yang, 2019) and in beauty-contests and generalized linear-quadratic games (e.g., Hellwig and Veldkamp 2009, Myatt and Wallace, 2012, Colombo et al, 2014, Pavan, 2016, and Hebert and La'O, 2020).⁴ More broadly, various papers consider information acquisition in Bayesian games with strategic complementarities (see, e.g., Lehrer and Rosenberg, 2006, and Amir and Lazzati, 2016, for some of the earlier references, and Banerjee et al. 2020 and Liang and Mu 2020 for recent developments). Information acquisition has also been examined in auctions (e.g., Persico 2000), in mechanism design (e.g., Bergemann and Välimäki, 2002), in contracting games (e.g., Crémer and Khalil, 1992, 1994, Crémer et al., 1998a,b, Dang 2008, Tirole 2009, Bolton and Faure-Grimaud 2010, Pavan and Tirole, 2021a,b), and in security design (e.g., Dang et al 2017, Farhi and Tirole 2015, and Yang, 2020). In the papers cited above, information acquisition is primarily about exogenous payoff states.⁵ Instead, Dewatripont and Tirole (2005), Che and Kartik (2009), Calvo-Armengol et al. (2015), Sethi and Yildiz (2016, 2018), Kozlovskaya (2018), and Adriani and Sonderegger (2020), and Denti (2020) consider the case of players learning about other players' beliefs, and/or communicating their view of the game to other players. Related to the literature on information acquisition is also the literature on rational inattention. See Sims (2003) for one of the earliest contributions, Maćkowiak and Wiederholt (2009), Matejka and McKay (2012), and Matejka et al (2017) for some of the recent contributions, and Maćkowiak et al. (2020) for an overview of this literature.⁶

A second strand is the literature on sparsity and endogenous depth of reasoning. See Gabaix (2014) for the former and Alaoui and Penta (2016a, 2016b, 2018) for the latter. Related is also the literature on psychological games (e.g., Geanakoplos et al., 1988, and Battigalli and Dufwenberg, 2009). This literature looks at games where players' (first and higher-order) beliefs over other players' intentions enter directly into payoffs, for example in the form of sequential reciprocity and regret. In contrast, in the present paper, we consider situations in which players learn, at a cost, about a game, but where the knowledge they acquire is instrumental to the selection of the best responses in the primitive game.

A third strand is the recent literature on persuasion and information design (see Bergemann

⁴See also Bergemann and Morris (2013) for an analysis of robust predictions in these games.

⁵Denti (2020) and Hebert and La'O (2020), however, also consider the case of players learning about other players' actions. See also Angeletos and Sastry (2020) for the validity of welfare theorems in inattentive economies where agents observe other agents' actions and prices with endogenous noise.

⁶Compared to the literature on information acquisition, the literature on rational inattention considers more flexible specifications of the signal technology but imposes specific assumptions on the cost functional, with the latter typically taking the form of entropy reduction. See Caplin et al. (2018), Hebert and La'O (2020), and Hebert and Woodford (2020) for related models of rational inattention with alternative cost functionals.

and Morris, 2019, and Kamenica, 2019, for overviews), as well as the literature on signal jamming and career concerns (see Fudenberg and Tirole, 1986, Holmström, 1999, and Dewatripont et al., 1999, for earlier contributions and Horner and Lambert, 2019, for recent developments). Examples of manipulative cognition can also be found in the literature on framing and memory management (see, e.g., Benabou and Tirole, 2002, Mullainathan, 2002, Dessi, 2008, Mullainathan et al. 2008, Benabou, 2013, Gottlieb 2014a,b, and Salant and Siegel, 2018).

The strands of the literature mentioned above are too vast to be successfully summarized here. While the analysis in the present paper touches upon themes considered in the aforementioned body of work, to the best of our knowledge, this is the first paper to introduce the notion of expectation conformity in strategic reasoning and to investigate the systematic role that synergies in cognition play for the determinacy of equilibria and the selection of the actual cognitive profiles.

2 Cognitive games, expectation conformity, and equilibrium determinacy

2.1 Cognitive games

Players, actions, and payoffs. There are $n \in \mathbb{N}$ players, indexed by $i \in I \equiv \{1, \dots, n\}$, with $n \geq 2$, engaged in a primitive normal- or extensive-form game. This primitive game will also be referred to as the “stage-2” or the “downstream” game, with the three expressions meant to be synonyms. In this stage-2 game, player i has action space A_i and receives a gross payoff $u_i(\alpha_i, \alpha_{-i}, \omega)$ in state of nature $\omega \in \Omega$, where $\alpha_i \in \Delta(A_i)$ is player i ’s mixed action and $\alpha_{-i} \equiv (\alpha_1, \dots, \alpha_{i-1}, \alpha_{i+1}, \dots, \alpha_n) \in \prod_{j \neq i} \Delta(A_j)$ is a profile of mixed actions for players other than i . When α_i is a Dirac measure assigning probability one to action $a_i \in A_i$ and likewise α_{-i} is a collection of Dirac measures with each α_j assigning probability one to some action $a_j \in A_j$, $j \neq i$, we abuse notation and denote by $u_i(a_i, a_{-i}, \omega)$ player i ’s payoff in state ω when the action profile is (a_i, a_{-i}) .

Cognition. As anticipated in the Introduction, cognition may take various forms. Here, we consider the case where it determines a collection of information structures, which is the case favored in the literature. In Section 6, we consider the case where cognition determines the players’ ability to compute iterated best responses, as in the level-k model.

The players have a common prior F on the state space Ω .⁷ Prior to playing the stage-2 game, the players *privately* engage in cognition at stage 1, resulting in a vector of information structures. We index player i ’s cognitive activity by the parameter ρ_i and then denote by $C_i(\rho_i)$ player i ’s cost of selecting cognition ρ_i .⁸

In some applications, only one of the players engages in cognition; this amounts to the other

⁷Che and Kartik (2009) by contrast look at incentives to acquire information in an environment with heterogeneous priors.

⁸This modeling of information structures applies not only to standard models of search and information acquisition, but also to endogenous imperfect recall, as well as to categorical thinking (e.g. Mullainathan 2002 and Mullainathan et al 2008). It does not apply though to the case of endogenous depth of reasoning in the level-k model, the description of which is postponed to Section 6.

players' having an infinite cost except for a "no cognition" benchmark; we will call this case "one-sided cognition".

Player i 's stage-2 action must be measurable with respect to the information structure induced by the stage-1 cognitive choices $\rho \equiv (\rho_i, \rho_{-i})$, with $\rho_{-i} \equiv (\rho_1, \dots, \rho_{i-1}, \rho_{i+1}, \dots, \rho_n)$.⁹ We capture the above measurability constraints as follows. Given ρ , at stage 2, each player i observes a signal realization $s_i \in S_i$ with the vector $s \equiv (s_1, \dots, s_n)$ drawn from $S \equiv \prod_{i=1}^n S_i$ according to the distribution $Q(s|\omega, \rho)$.

We say that cognition is *self-directed* when the information that each player receives is independent of the other players' cognitive choices. That is, there exists a collection of distributions $(Q_i)_{i=1}^n$ such that, for any (s, ω, ρ) , $Q(s|\omega, \rho) = \prod_{i \in I} Q_i(s_i|\omega, \rho_i)$. A special case of self-directed cognition that is prominent in applications is when the information structures are ordered. Player i 's choice of information structure is then represented by $\rho_i \in \mathbb{R}$ and, for $\rho_i < \rho'_i$, player i ' signal under ρ'_i is Blackwell more informative than under ρ_i .¹⁰ It is natural in most applications to assume that, in this case, C_i is monotonically increasing in ρ_i : A more informative signal is cognitively more expensive.¹¹

We also consider the case of *manipulative* cognition, that is, cognitive activities through which a player affects the information of other players. Examples of manipulative cognition include signal jamming, framing, and information disclosure. In this case, irrespectively of whether or not the signals s are drawn independently across the players, conditional on ω , the distribution from which a player's signal is drawn depends on other players' cognitive choices ρ_{-i} .

Strategies, reduced-form payoffs, and equilibria. A (mixed) strategy $\sigma_i \in \Delta(A_i)^{S_i}$ for player i in the continuation game that starts after players' cognition has been selected is a mapping $\sigma_i : S_i \rightarrow \Delta(A_i)$ that specifies, for each signal realization $s_i \in S_i$, a mixed action $\sigma_i(s_i) \in \Delta(A_i)$.

Given cognition ρ and continuation strategy profile σ , player i 's gross expected payoff is equal to

$$U_i(\sigma; \rho) \equiv \int_{\omega} \left[\int_s u_i(\sigma_i(s_i), \sigma_{-i}(s_{-i}), \omega) dQ(s|\omega, \rho) \right] dF(\omega)$$

Player i 's *net* payoff is equal to the above gross payoff, minus the stage-1 cognitive cost $C_i(\rho_i)$.

For any vector of cognitive choices ρ , let σ^ρ denote the associated stage-2 continuation equilibrium strategy profile. We assume that, for any ρ , either the stage-2 continuation equilibrium is unique, or some selection is in place.¹²

Next, for any ρ and ρ'_i , let

⁹Messages and disclosure decisions, if any, are part of the stage-2 strategies in this formulation.

¹⁰When signals take the form of partitions of Ω , $\rho_i < \rho'_i$ means that ρ'_i is finer than ρ_i .

¹¹This need not be the case for all applications though. Consider strategic memory management: Increasing the probability of forgetting some information that one has received (repression) is likely to be costly. By contrast, the case in which a player receives two pieces of information simultaneously when searching and would have to pay an extra cost to receive only one (unbundling) is not problematic in our interpersonal *covert*-information-acquisition context: the unbundled information structure is simply irrelevant and can be assumed to be infinitely costly (this would not be the case with *overt* information acquisition since a player may suffer when other players know that he has more information).

¹²Existence of stage-2 equilibria follows from standard assumptions.

$$V_i(\rho'_i; \rho) \equiv \sup_{\sigma_i \in \Delta(A_i)^{S_i}} U_i(\sigma_i, \sigma_{-i}^\rho; \rho'_i, \rho_{-i})$$

denote the maximal gross payoff that player i can obtain by covertly selecting cognition ρ'_i and then adjusting his stage-2 behavior optimally, when all players other than i play according to σ_{-i}^ρ . For any ρ ,

$$V_i(\rho_i; \rho) = U_i(\sigma^\rho; \rho)$$

then denotes the gross payoff that player i obtains by selecting the same cognition ρ_i as specified in the profile $\rho = (\rho_i, \rho_{-i})$ and then playing according to the continuation equilibrium σ^ρ .

When cognition is self-directed, and when different cognitive levels are ranked so that, for any $\rho'_i > \rho_i$, the signal player i receives under ρ'_i is Blackwell more informative than the one he receives under ρ_i , then $V_i(\rho'_i; \rho) \geq V_i(\rho_i; \rho)$ for any $\rho'_i > \rho_i$ (more information never hurts, provided that it is self-directed and that σ_{-i} is held fixed at σ_{-i}^ρ). However, $V_i(\rho'_i; (\rho'_i, \rho_{-i}))$ can be smaller than $V_i(\rho_i; (\rho_i, \rho_{-i}))$: As is well-known, a player may suffer from being more informed when the other players adjust their behavior in response to the expectation that player i is more informed, that is, if they play according to $\sigma_{-i}^{\rho'_i, \rho_{-i}}$ instead of $\sigma_{-i}^{\rho_i, \rho_{-i}}$.

Hereafter, for simplicity, we will be focusing on equilibria in which players do not mix over their cognitive choice.¹³ In equilibrium, all players have correct expectations about the cognitive choices of all other players, even though they do not directly observe the actual cognitive choices. We will say that the cognitive profile ρ is an “equilibrium” if, for any player i and any ρ'_i ,

$$V_i(\rho_i; \rho) - C_i(\rho_i) \geq V_i(\rho'_i; \rho) - C_i(\rho'_i).$$

2.2 Expectation conformity

Let $\rho = (\rho_i, \rho_{-i})$ and $\hat{\rho} = (\hat{\rho}_i, \hat{\rho}_{-i})$ denote two arbitrary cognitive profiles.

Definition 1 (expectation conformity). Expectation conformity (EC) holds for cognitive profiles ρ and $\hat{\rho}$ if, for all i ,¹⁴

$$V_i(\hat{\rho}_i; \hat{\rho}) - V_i(\rho_i; \hat{\rho}) \geq V_i(\hat{\rho}_i; \rho) - V_i(\rho_i; \rho). \quad \left(EC_{\{\rho, \hat{\rho}\}} \right)$$

Within and across complementarities.

¹³If stage 2 corresponds to an extensive-form game and player i mixes over his cognition, then player i 's early actions in the stage-2 game might reveal something about her actual choice of cognition. This possibility is interesting, but naturally brings multiplicity driven by out-of-equilibrium beliefs. We do not entertain it here.

¹⁴Note the importance of covert investments for this condition. Were investments in information overt, the condition would become

$$V_i(\hat{\rho}_i; \hat{\rho}_i, \hat{\rho}_{-i}) - V_i(\rho_i; \rho_i, \hat{\rho}_{-i}) \geq V_i(\hat{\rho}_i; \hat{\rho}_i, \rho_{-i}) - V_i(\rho_i; \rho_i, \rho_{-i}).$$

That is, expectation conformity reflects the fact that the players do not observe each other's choice of information structure, whereas the above condition posits that information structures are common knowledge at stage 2 *on and off* the equilibrium path.

Let us decompose the difference

$$\Gamma_i^{EC}(\rho, \hat{\rho}) \equiv \left[V_i(\hat{\rho}_i; \hat{\rho}_i, \hat{\rho}_{-i}) - V_i(\rho_i; \hat{\rho}_i, \hat{\rho}_{-i}) \right] - \left[V_i(\hat{\rho}_i; \rho_i, \rho_{-i}) - V_i(\rho_i; \rho_i, \rho_{-i}) \right]$$

between the left-hand and the right-hand side of $EC_{\{\rho, \hat{\rho}\}}$ into a within-player *unilateral expectation conformity (UEC)* term

$$\Gamma_i^{UEC}(\rho, \hat{\rho}) \equiv \left[V_i(\hat{\rho}_i; \hat{\rho}_i, \rho_{-i}) - V_i(\rho_i; \hat{\rho}_i, \rho_{-i}) \right] - \left[V_i(\hat{\rho}_i; \rho_i, \rho_{-i}) - V_i(\rho_i; \rho_i, \rho_{-i}) \right]$$

which captures the impact of the other players' anticipation of player i 's cognition, fixing the other players' levels of cognition ρ_{-i} , and an across-players *increasing differences (ID)* term,

$$\Gamma_i^{ID}(\rho, \hat{\rho}) \equiv \left[V_i(\hat{\rho}_i; \hat{\rho}_i, \hat{\rho}_{-i}) - V_i(\rho_i; \hat{\rho}_i, \hat{\rho}_{-i}) \right] - \left[V_i(\hat{\rho}_i; \hat{\rho}_i, \rho_{-i}) - V_i(\rho_i; \hat{\rho}_i, \rho_{-i}) \right]$$

which by contrast captures the impact of the other players' own cognition.

When expectation conformity holds for cognitive profiles ρ and $\hat{\rho}$, i.e., when $\Gamma_i^{EC}(\rho, \hat{\rho}) \geq 0$, it is interesting to investigate whether this comes from unilateral expectation conformity, i.e., $\Gamma_i^{UEC}(\rho, \hat{\rho}) \geq 0$, from variations in the other players' cognition, i.e., $\Gamma_i^{ID}(\rho, \hat{\rho}) \geq 0$, or from both. Clearly, in games in which only one player engages in cognition (*one-sided cognitive games*), because, for all $(\rho, \hat{\rho})$, necessarily $\rho_{-i} = \hat{\rho}_{-i}$ (players other than i have only one feasible cognitive choice), then $\Gamma_i^{ID}(\rho, \hat{\rho}) = 0$ and hence

$$\Gamma_i^{EC}(\rho, \hat{\rho}) = \Gamma_i^{UEC}(\rho, \hat{\rho}).$$

To illustrate the possibility that EC arises from ID rather than from UEC, consider a simple *matching model* in which two players may invest in recognizing what's in it for them in a given partnership; that is, each potential match is characterized by a positive surplus for player i if the partner is adequate and a highly negative payoff otherwise. Hence, a match occurs only if both players can ascertain it is a good one for them. More formally, in this example, an information structure for player i coincides with the probability $\rho_i \in [0, 1]$ that player i perfectly learns the quality of the match for him (at some cost $C_i(\rho_i)$). With the complementary probability, player i learns nothing. At stage 2, the two players each have a veto right on the matching. That is, matching occurs if and only if both players approve it. Eliminating weakly dominated strategies, each player's stage-2 behavior is then independent of her expectation about the other player's cognition. A player who knows the match to be of high quality for herself approves the match. A player who, instead, is either uninformed, or knows the match to be of low quality for herself, rejects the match. In this game, for any $\rho = (\rho_i, \rho_j)$ and $\hat{\rho} = (\hat{\rho}_i, \hat{\rho}_j)$, $\Gamma_i^{UEC}(\rho, \hat{\rho}) = [\hat{\rho}_i \rho_j - \rho_i \rho_j] - [\hat{\rho}_i \rho_j - \rho_i \rho_j] = 0$. By contrast, $\Gamma_i^{ID}(\rho, \hat{\rho}) = (\hat{\rho}_i - \rho_i)(\hat{\rho}_j - \rho_j) > 0$, capturing the standard strategic complementarities that are conducive to equilibrium multiplicity at the cognition stage.

2.3 Equilibrium determinacy and cognitive traps

Revealed preferences imply that a necessary condition for ρ and $\hat{\rho}$ to be two equilibrium cognitive profiles (for some cost structures $(C_i(\cdot))_{i \in I}$) is that $\Gamma_i^{EC}(\rho, \hat{\rho}) \geq 0$ for all i . The following straightforward proposition says that this condition is also sufficient for ρ and $\hat{\rho}$ to be equilibrium profiles, for appropriately chosen cost functions:

Proposition 1 (equilibrium determinacy). (a) If $EC_{\{\rho, \hat{\rho}\}}$ is satisfied for two distinct cognitive profiles ρ and $\hat{\rho}$, then there exist cost functions $(C_i(\cdot))_{i \in I}$ such that both ρ and $\hat{\rho}$ are equilibrium cognitive profiles. Furthermore, if cognition is self-directed and totally ordered, and $\hat{\rho}_i$ is Blackwell more informative than ρ_i , for all i , the cost functions can be chosen to be monotonic. (b) If $EC_{\{\rho, \hat{\rho}\}}$ is not satisfied for any two distinct cognitive profiles $(\rho, \hat{\rho})$, then, irrespective of the cost functions $(C_i(\cdot))_{i \in I}$, the equilibrium is unique.

Proof. First, consider part (a). Suppose that $EC_{\{\rho, \hat{\rho}\}}$ is satisfied for the two distinct cognitive profiles $\rho = (\rho_i, \rho_{-i})$ and $\hat{\rho} = (\hat{\rho}_i, \hat{\rho}_{-i})$. For ρ and $\hat{\rho}$ to be equilibrium profiles, it must be that, for all i

$$V_i(\hat{\rho}_i; \hat{\rho}) - V_i(\rho_i; \hat{\rho}) \geq C_i(\hat{\rho}_i) - C_i(\rho_i) \geq V_i(\hat{\rho}_i; \rho) - V_i(\rho_i; \rho). \quad (1)$$

It then suffices to pick cost functions that, in addition to (1), satisfy $C_i(\tilde{\rho}_i) = +\infty$ if $\tilde{\rho}_i \notin \{\rho_i, \hat{\rho}_i\}$. When cognition is self-directed and totally ordered, and $\rho_i \leq \hat{\rho}_i$ for all i , meaning that $\hat{\rho}_i$ is Blackwell more informative than ρ_i , all i , let $(C_i(\cdot))_{i \in I}$ be any profile of cost functions that, in addition to (1), satisfy

$$C_i(\tilde{\rho}_i) = \begin{cases} C_i(\rho_i) & \text{for } \tilde{\rho}_i \leq \rho_i \\ C_i(\hat{\rho}_i) & \text{for } \rho_i < \tilde{\rho}_i \leq \hat{\rho}_i \\ +\infty & \text{for } \tilde{\rho}_i > \hat{\rho}_i. \end{cases}$$

Because more information cannot hurt a player when covertly acquired, then $V_i(\hat{\rho}_i; \rho) - V_i(\rho_i; \rho) \geq 0$.¹⁵ Hence, $C_i(\hat{\rho}_i) \geq C_i(\rho_i)$, implying that the above cost functions are monotone. It is then immediate to see that ρ and $\hat{\rho}$ are indeed equilibrium profiles for these cost functions.

Next, consider part (b). If ρ and $\hat{\rho}$ are two distinct equilibria, Condition (1) must be satisfied, and so $EC_{\{\rho, \hat{\rho}\}}$ must hold. Hence if EC holds for no pair of cognitive profiles, then no matter the cost functions $(C_i(\cdot))_{i \in I}$, the equilibrium is unique. \square

Definition 2 (cognitive traps). Suppose that cognition is self-directed and totally ordered. Players are exposed to a cognitive trap if there exist two equilibrium cognitive profiles ρ and $\hat{\rho}$ such that, for all i for whom $\hat{\rho}_i \neq \rho_i$, (a) $\hat{\rho}_i$ is Blackwell more informative than ρ_i , and (b) $V_i(\hat{\rho}_i; \hat{\rho}) - C_i(\hat{\rho}_i) < V_i(\rho_i; \rho) - C_i(\rho_i)$.

2.4 Local expectation conformity and strategic complementarity/substitutability in cognition

Suppose that there are two players, that cognition is totally ordered, that $\rho_i \in \mathbb{R}$, and that, for given cost functions, there is a unique and stable equilibrium cognitive profile, denoted by ρ , with $\rho_i > 0$, all i . Further assume that $V_i(\rho'_i; \hat{\rho})$ are twice continuously differentiable, and that, for any $\hat{\rho}$, $V_i(\cdot; \hat{\rho}) - C_i(\cdot)$ is concave in ρ'_i .

The following definition extends the notion of unilateral expectation conformity to local changes in cognition around the equilibrium level:

¹⁵This property, however, need not hold when players are boundedly rational; see the discussion in Section 6.

Definition 3. Positive (alternatively, negative) *local UEC* holds at the equilibrium cognitive profile $\rho = (\rho_1, \rho_2)$ if

$$\left. \frac{\partial^2 V_i}{\partial \rho'_i \partial \rho_i}(\rho'_i; \rho_i, \rho_j) \right|_{\rho'_i = \rho_i} > 0 \quad (2)$$

(alternatively, if the above inequality is reversed) for $i, j = 1, 2, j \neq i$.

Likewise, the following definition extends the notion of increasing differences to local changes in cognition around the equilibrium level:

Definition 4. Positive (alternatively, negative) *local ID* holds at the equilibrium cognitive profile $\rho = (\rho_1, \rho_2)$ if

$$\left. \frac{\partial^2 V_i}{\partial \rho'_i \partial \rho_j}(\rho'_i; \rho_i, \rho_j) \right|_{\rho'_i = \rho_i} > 0, \quad (3)$$

(alternatively, if the above inequality is reversed) for $i, j = 1, 2, j \neq i$.

Local expectation conformity is then defined as follows:

Definition 5. Positive (alternatively, negative) *local EC* holds at the equilibrium cognitive profile $\rho = (\rho_1, \rho_2)$ if

$$\left. \frac{\partial^2 V_i}{\partial \rho'_i \partial \rho_i}(\rho'_i; \rho_i, \rho_j) \right|_{\rho'_i = \rho_i} + \left. \frac{\partial^2 V_i}{\partial \rho'_i \partial \rho_j}(\rho'_i; \rho_i, \rho_j) \right|_{\rho'_i = \rho_i} > 0 \quad (4)$$

(alternatively, if the above inequality is reversed) for $i, j = 1, 2, j \neq i$.

Now let

$$R_i(\rho) \equiv \arg \max_{\rho'_i} \{V_i(\rho'_i; \rho) - C_i(\rho'_i)\}$$

denote player i 's *cognitive reaction curve*.

Definition 6. Cognitive choices are local strategic complements (alternatively, substitutes) at the equilibrium cognitive profile $\rho = (\rho_1, \rho_2)$ if $R_i(\rho)$ is increasing (alternatively, decreasing) in ρ_j , $i, j = 1, 2, j \neq i$.

The following proposition illustrates that ID and UEC play a role not only for the determinacy of equilibria but also for the strategic complementarity/substitutability in cognition.

Proposition 2. *Suppose that there are two players, that the payoff functions $V_i(\rho'_i; \hat{\rho})$ are twice continuously differentiable, and that, for any $\hat{\rho}$, $V_i(\cdot; \hat{\rho})$ is concave in ρ'_i , $i = 1, 2$. Suppose that negative local UEC holds at the equilibrium cognitive profile $\rho = (\rho_1, \rho_2)$. Then cognitive choices are local strategic complements (alternatively, local strategic substitutes) at $\rho = (\rho_1, \rho_2)$ if positive (alternatively, negative) local ID holds. When, instead, positive local UEC holds and*

$$\left. \frac{\partial^2 V_i}{\partial \rho'_i \partial \rho_i}(\rho'_i; \rho_i, \rho_j) \right|_{\rho'_i = \rho_i} > C_i''(\rho_i) - \left. \frac{\partial^2 V_i}{\partial (\rho'_i)^2}(\rho'_i; \rho_i, \rho_j) \right|_{\rho'_i = \rho_i}$$

then cognitive choices are local strategic complements (alternatively, local strategic substitutes) at the equilibrium cognitive profile $\rho = (\rho_1, \rho_2)$ if negative (alternatively, positive) local ID holds.

Proof. The equilibrium cognitive profile ρ must satisfy the first-order conditions

$$\left. \frac{\partial V_i}{\partial \rho'_i}(\rho'_i; \rho_i, \rho_j) \right|_{\rho'_i = \rho_i} - C'_i(\rho_i) = 0,$$

$i = 1, 2$. From the implicit function theorem,

$$R'_i(\rho_j) = \frac{- \left. \frac{\partial^2 V_i}{\partial \rho'_i \partial \rho_j}(\rho'_i; \rho_i, \rho_j) \right|_{\rho'_i = \rho_i}}{\left[\left. \frac{\partial^2 V_i}{\partial (\rho'_i)^2}(\rho'_i; \rho_i, \rho_j) \right|_{\rho'_i = \rho_i} - C''_i(\rho_i) \right] + \left. \frac{\partial^2 V_i}{\partial \rho'_i \partial \rho_i}(\rho'_i; \rho_i, \rho_j) \right|_{\rho'_i = \rho_i}}.$$

The term in square brackets is negative by the concavity of $V_i(\cdot; \rho) - C_i(\cdot)$. The results then follow from the fact that the sign of the numerator is determined by ID and the sign of the second term in the denominator is the sign of UEC. \square

3 Constant-sum games

Before we move on to analyze classes of games that satisfy expectation conformity, at least for certain cognitive profiles, it is interesting to consider an important class that does not satisfy it (at least not strictly), no matter the profiles. Suppose that the stage-2 game is a constant-sum game between two players. That is, for all $(\alpha_i, \alpha_j, \omega)$, the gross payoffs satisfy the constant-sum condition:

$$u_i(\alpha_i, \alpha_j, \omega) + u_j(\alpha_i, \alpha_j, \omega) = k(\omega),$$

where k is an arbitrary function of the state of Nature. The overall game obviously is not a zero-sum game: Any cognition, when costly, necessarily reduces total surplus and amounts to pure rent-seeking.

Constant-sum games have several remarkable properties. For example, a player can only benefit from having (and being known to have) more information (Lehrer and Rosenberg 2006), a property that is well known to be violated for general games. Another interesting property is given by the following result:¹⁶

Proposition 3 (constant-sum games). *Two-person constant-sum games satisfy the following property, for all $(\rho, \hat{\rho})$:*

$$\Sigma_i \Gamma_i^{EC}(\rho, \hat{\rho}) \leq 0.$$

As a consequence, if there are multiple equilibria, in none of them can a player have a strict preference for her equilibrium cognition over her cognition in any other equilibrium.

Proof. The constant-sum property implies that, for any pair of cognitive profiles $(\rho, \hat{\rho})$,

$$\Sigma_i \left\{ \left[U_i(\sigma_i^{\hat{\rho}}, \sigma_j^{\hat{\rho}}; \hat{\rho}_i, \hat{\rho}_j) - U_i(\sigma_i^{\rho}, \sigma_j^{\hat{\rho}}; \rho_i, \hat{\rho}_j) \right] - \left[U_i(\sigma_i^{\hat{\rho}}, \sigma_j^{\rho}; \hat{\rho}_i, \rho_j) - U_i(\sigma_i^{\rho}, \sigma_j^{\rho}; \rho_i, \rho_j) \right] \right\} = 0. \quad (5)$$

Next observe that, for any $l, m \in \{1, 2\}$, $m \neq l$,

$$U_l(\sigma_l^{\rho}, \sigma_m^{\hat{\rho}}; \rho_l, \hat{\rho}_m) \leq \sup_{\sigma_l \in \Delta(A_l)^{S_l}} U_l(\sigma_l, \sigma_m^{\hat{\rho}}; \rho_l, \hat{\rho}_m)$$

¹⁶We are grateful to Gabriel Carroll for conjecturing that zero-sum games fail to satisfy expectation conformity.

and

$$U_l(\sigma_l^{\hat{\rho}}, \sigma_m^{\rho}; \hat{\rho}_l, \rho_m) \leq \sup_{\sigma_l \in \Delta(A_l)^{S_l}} U_l(\sigma_l, \sigma_m^{\rho}; \hat{\rho}_l, \rho_m).$$

Hence, Condition (5) implies that

$$\Sigma_i \Gamma_i^{EC}(\rho, \hat{\rho}) \leq 0$$

for all $(\rho, \hat{\rho})$. Because equilibrium multiplicity requires $\Gamma_i^{EC}(\rho, \hat{\rho}) \geq 0$ for all i , the above inequality implies that, if there are multiple equilibria, then $\Gamma_i^{EC}(\rho, \hat{\rho}) = 0$ for all i . Using ((1)), we then have that, if ρ and $\hat{\rho}$ are both equilibrium cognitive profiles, then necessarily

$$V_i(\hat{\rho}_i; \hat{\rho}) - V_i(\rho_i; \hat{\rho}) = C_i(\hat{\rho}_i) - C_i(\rho_i) = V_i(\hat{\rho}_i; \rho) - V_i(\rho_i; \rho). \quad (6)$$

In each equilibrium, each player must thus be indifferent between selecting the cognitive level she is supposed to select in that equilibrium and the cognitive level she is supposed to selected in any other equilibrium. \square

That, in case of multiple equilibria, in each equilibrium, each player is indifferent between selecting her equilibrium cognitive level and the cognitive level she is supposed to select under any of the other equilibria does not mean that the equilibrium payoffs are the same in all equilibria. To illustrate, consider the zero-sum game in which i 's payoff is $(a_i - a_j)\omega$, for $i, j = 1, 2$ and $j \neq i$. Further assume that $A_m = \{1, -1\}$, for $m = 1, 2$, and that ω is drawn from $\Omega = \{-1, 1\}$ with equal probability. Lastly, suppose that $\rho_i \in \{1, \emptyset\}$, $i = 1, 2$. When $\rho_i = 1$, player i perfectly learns ω , that is, she receives a signal $s_i \in \{-1, 1\}$ with probability $\Pr(1|1) = \Pr(-1|-1) = 1$ that is perfectly informative of the state. When, instead, $\rho_i = 0$, player i receives no information about ω . Let $C_i(1) = 1$ and $C_i(\emptyset) = 0$. The game admits two pure-strategy equilibria. In the first one, both players learn the state and then match the state with their action. In the second equilibrium, neither player learns the state and then selects action $a_m = 1$ with probability one, $m = 1, 2$. The equilibrium payoffs under the first equilibrium is -1 whereas the equilibrium payoffs under the second equilibrium are equal to 0. Hence, in this game, the high-cognition equilibrium is a cognitive trap. It is also easy to see that, in each equilibrium, each player is indifferent between selecting her equilibrium cognition and the cognition specified in the other equilibrium.

Importantly, note that, in contrast to 2-player constant-sum games, n -player constant-sum games in which $n > 2$ may admit multiple strict equilibria. To see this, take a non-constant-sum two-player game admitting multiple strict equilibria (such as the coordination games studied in the next section). Have this game played twice, by players 1 and 2 and by players 3 and 4, respectively. Use players 3 and 4 (alternatively, players 1 and 2) as passive ‘‘budget balancers’’ in the game played by 1 and 2 (alternatively, 3 and 4). The transformed game is a constant-sum game that admits multiple strict equilibria.

4 Self-directed cognition

In this section, we specialize the analysis to games in which cognition is self-directed, meaning that each player's cognition is unaffected by the other players' cognitive posture. First, we examine the case where cognition takes the form of sparsity and then the case where it takes the more familiar

form of noisy information acquisition. We conclude by considering the case where cognition takes the form of “spying” on other players’ signals.

Throughout the analysis, we assume that players’ payoff are given by

$$u_i(a_i, a_{-i}, \omega) = -(1 - \beta)(a_i - g(\omega))^2 - \beta(a_i - \bar{a}_{-i})^2 + \psi(a_{-i}, \omega), \quad (7)$$

where $a_i \in A_i = \mathbb{R}$, $a_{-i} \equiv (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n) \in \mathbb{R}^{n-1}$, $\omega \equiv (\omega^k)_{k=1}^{k=K} \in \mathbb{R}^K$, for some $K \in \mathbb{N} \cup \{+\infty\}$ $\bar{a}_{-i} \equiv \sum_{j \neq i} a_j / (n - 1) \in \mathbb{R}$, $g : \mathbb{R}^K \rightarrow \mathbb{R}$, and $\psi : \mathbb{R}^{K+n-1} \rightarrow \mathbb{R}$. The payoff state ω is thus a collection of “fundamental variables.” The function g aggregates such variables into a uni-dimensional statistics $g(\omega)$. The variable $\bar{a}_{-i} \equiv \sum_{j \neq i} a_j / (n - 1)$ is the average action of player i ’s opponents, and the scalar $\beta \in \mathbb{R}$ parametrizes the intensity of the strategic interactions, with $0 < |\beta| < 1$. The case $\beta > 0$ corresponds to a strategic situation in which actions are *strategic complements*, whereas the case $\beta < 0$ corresponds to a situation in which actions are *strategic substitutes*.¹⁷ Finally, the function $\psi(a_{-i}, \omega)$ summarizes various external effects that matter for payoffs but do not play any role for the selection of the individual best responses.

4.1 Sparsity in games

Consider a situation in which cognition takes the form of a player choosing which dimensions of the state ω to learn about. As in Gabaix (2014), each dimension that is not learnt is treated as if it did not exist. While sparsity is typically interpreted as a form of bounded rationality, below we introduce a formalization that captures the spirit of sparsity, while retaining the convenience of fully rational behavior.

The statistics $g(\omega)$ takes the form $g(\omega) = (1 + \sum_{k=1}^K \omega^k) / (1 - \beta)$. Such specification of the g function is not essential for the results but facilitates certain derivations. It is commonly believed that each dimension ω^k is drawn independently from the other dimensions from a distribution F^k with zero mean and variance σ_k^2 .

There is a natural progression in reasoning: each player learns the realization of the various dimensions of ω in sequence. That is, learning the realization of ω^k requires having learnt the first $k - 1$ dimensions $(\omega^1, \dots, \omega^{k-1})$. In this environment, a player’s cognition $\rho_i \in \mathbb{N}$ thus corresponds to her “depth of knowledge,” that is, to the number of dimensions of the state ω the player learns about.

The above formalization captures the distinctive feature of sparsity according to which a player who decides to learn ρ_i dimensions of the state reasons “as if” the remaining $k - \rho_i$ dimensions did not exist (that is, as if the state had only ρ_i dimensions). A second feature of the sparsity model that is related to the property discussed above is that a player who goes deeper in the exploration of the state is able to perfectly predict the opponent’s behavior, whereas a player who explores fewer dimensions than the opponent reasons (and acts) as if the opponent explored the same dimensions that she explored, even if she knows that this is not the case.

To see why the above properties hold, but also to appreciate their implications for expectation conformity and equilibrium determinacy, let $\omega_i \equiv (\omega^1, \dots, \omega^{\rho_i})$ be the subset of the state ω explored

¹⁷The formulas below extend to the case in which $\beta = 0$, but this case is not interesting, for it corresponds to a non-strategic situation.

by player i . Using the general notation of Section (2.1), we then have that, for any (ω, ρ) , $Q(s|\omega, \rho)$ is a Dirac measure assigning probability one to the signal vector $s = (\omega_1, \dots, \omega_n)$, where n is the number of players.

For simplicity, assume that $n = 2$ and, without loss of generality, then let player 1 be the player with the lowest depth of knowledge, that is, for whom $\rho_1 \leq \rho_2$.¹⁸ Given the cognitive profile $\rho = (\rho_1, \rho_2)$, one can show that, in the stage-2 game, there exists a unique continuation equilibrium and is such that, for any ω , the two players' equilibrium strategies are Dirac distributions assigning probability one to the actions¹⁹

$$a_1^\rho(\omega_1) = \frac{1 + \sum_{k=1}^{\rho_1} \omega^k}{1 - \beta}$$

and²⁰

$$a_2^\rho(\omega_2) = \frac{1 + \sum_{k=1}^{\rho_1} \omega^k}{1 - \beta} + \sum_{k=\rho_1+1}^{\rho_2} \omega^k$$

It is then easy to verify that each player reasons about the state and acts as if each neglected dimension $k > \rho_i$ did not exist. In particular, the equilibrium action of the player who is known to be behind in the exploration of the state (player 1) is invariant in how far ahead the opponent is in the exploration of the state and coincides with the equilibrium action $a_1^{\rho_1, \rho_1}(\omega_1)$ that player 1 would choose if the state was commonly known to have only ρ_1 dimensions and if all dimensions $\omega_1 = (\omega^1, \dots, \omega^{\rho_1})$ were common known to both players. Similarly, the stage-2 equilibrium action for the player who is known to be ahead in the exploration of the state (player 2) coincides with the equilibrium action of the player who is behind, $a_1^\rho(\omega_1)$, augmented by the extra knowledge $\sum_{k=\rho_1+1}^{\rho_2} \omega^k$ that player 2 has about the gross return to her action. Again, $a_2^\rho(\omega_2)$ also coincides with the player's equilibrium action when she expects the state to have only ρ_2 dimensions and when she expects her opponent to be aware of only the first $\rho_1 < \rho_2$ dimensions.

Given cognition $\rho = (\rho_1, \rho_2)$, the two players' ex-ante expected gross payoffs (up to a normalization eliminating all external effects that have no impact on individual best responses) when it is common knowledge that the two players engaged in cognition ρ , are then equal to

$$V_1(\rho_1; \rho) = \frac{1 + \sum_{k=1}^{\rho_1} \sigma_k^2}{(1 - \beta)^2}$$

and

$$V_2(\rho_2; \rho) = V_1(\rho_1; \rho) + \sum_{k=\rho_1+1}^{\rho_2} \sigma_k^2.$$

¹⁸The qualitative insights carry over to the case $n > 2$. The structure of the best responses, however, is more tedious than in the $n = 2$ case. When ρ and $\hat{\rho}$ are symmetric profiles (i.e., $\rho_i = \rho_j$ and $\hat{\rho}_i = \hat{\rho}_j$ for all $i, j \in I$), the results nor the $n > 2$ case coincide with those for the $n = 2$ case.

¹⁹The computation of the equilibrium actions, as well as the derivation of the interim expected gross payoffs below, is in the Appendix in the proof of Proposition (4).

²⁰Given the linear-quadratic specification in (7), and the fact that, in this application, $g(\omega) = (1 + \bar{\omega})/(1 - \beta)$, one can think of dimension $\omega^0 = 1$ as the part of the gross return to "investment" that is exogenously known to both players, $\omega^k/(1 - \beta)$ as the contribution of dimension $k = 0, 1, \dots$ to such gross return when commonly known, and ω^k as the contribution of dimension $k = 1, 2, \dots$ when individually known.

Next suppose that the two players are expected to engage in cognition $\rho = (\rho_1, \rho_2)$ and consider the effects of possible deviations by either player. Consider first player 1. In the Appendix, we show that if player 1 were to deviate to $\rho'_1 < \rho_1$, thus exploring fewer dimensions than expected to, then for any $\omega'_1 \equiv (\omega^1, \dots, \omega^{\rho'_1})$, her optimal stage-2 action

$$\frac{1 + \sum_{k=1}^{\rho'_1} \omega^k}{1 - \beta}$$

would coincide with the equilibrium action $a_1^{\rho'_1, \rho_2}(\omega_1)$ that she would select if the deviation was observable. This is because, after the deviation to $\rho'_1 < \rho_1$, player 1 reasons as if the state had only ρ'_1 dimensions and all such dimensions were known also to player 2, exactly as when the deviation is observable. Clearly, if the deviation was observable, player 1 would expect player 2 to change her response to the dimensions $(\rho'_1 + 1, \dots, \rho_1)$, adjusting to the fact that such dimensions are no longer commonly learnt. However, because player 1 no longer possesses any information about such dimensions once she deviates to $\rho'_1 < \rho_1$, the adjustment in player 2's action that would occur if the deviation was observable would not affect player 1's optimal action. Player 1's ex-ante gross expected payoff following the deviation to $\rho'_1 < \rho_1$ is thus equal to $V_1(\rho'_1; \rho) = \left[1 + \sum_{k=1}^{\rho'_1} \sigma_k^2\right] / (1 - \beta)^2$ and coincides with her equilibrium payoff $V_1(\rho'_1; (\rho'_1, \rho_2))$ under the profile (ρ'_1, ρ_2) .²¹

When, instead, player 1 deviates to cognition $\rho_1 < \rho'_1 \leq \rho_2$, that is, when she expands her depth of knowledge vis-a-vis what she is expected to learn, without, however, overtaking the rival, her optimal stage-2 action is then equal to

$$a_1^{\rho_1, \rho_2}(\omega_1) + (1 + \beta) \sum_{k=\rho_1+1}^{\rho'_1} \omega^k$$

and coincides with her stage-2 equilibrium action $a_1^{\rho_1, \rho_2}(\omega_1) = \left[1 + \sum_{k=1}^{\rho_1} \omega^k\right] / (1 - \beta)$ in the absence of the deviation, augmented by a term $(1 + \beta) \sum_{k=\rho_1+1}^{\rho'_1} \omega^k$ that reflects the extra return originating from both players knowing (and responding to) the dimensions $(\rho_1 + 1, \dots, \rho'_1)$, without, however, player 2 adjusting her response to account for the fact that such dimensions are now commonly known. Player 1's ex-ante gross expected payoff following such a deviation is then equal to

$$V_1(\rho'_1; \rho) = V_1(\rho_1; \rho) + (1 + \beta)^2 \sum_{k=\rho_1+1}^{\rho'_1} \sigma_k^2.$$

Finally, consider deviations to a cognitive level $\rho'_1 > \rho_2$. In this case, player-1's stage-2 optimal action is equal to

$$a_1^{\rho_1, \rho_2}(\omega_1) + (1 + \beta) \sum_{k=\rho_1+1}^{\rho_2} \omega^k + \sum_{k=\rho_2+1}^{\rho'_1} \omega^k \quad (8)$$

where, relative to the case $\rho_1 < \rho'_1 \leq \rho_2$ discussed above, the new term in (8) accounts for the fact that player 1 is the sole learner of the dimensions $(\rho_2 + 1, \dots, \rho'_1)$ and hence the return to learning such dimensions is $\sum_{k=\rho_2+1}^{\rho'_1} \omega^k$. Her ex-ante gross expected payoff following the deviation to $\rho'_1 > \rho_2$ is then equal to

$$V_1(\rho'_1; \rho) = V_1(\rho_1; \rho) + (1 + \beta)^2 \sum_{k=\rho_1+1}^{\rho_2} \sigma_k^2 + \sum_{k=\rho_2+1}^{\rho'_1} \sigma_k^2.$$

²¹Which also coincides with the payoff $V_1(\rho'_1; (\rho'_1, \rho'_1))$ when player 2 explores the same number of dimensions $\rho_2 = \rho'_1$ as player 1 (which also coincides with the equilibrium payoff when the state has only ρ'_1 dimensions, that is when $K = \rho'_1$). As anticipated above, these properties are distinctive features of sparsity in cognitive games.

Next, consider deviations by player 2. Because player 1's estimate of the (sum of the) dimensions for which player 2 is the sole learner is invariant to the number of dimensions that are learned solely by player 2, another feature of the sparsity model is that, if player 2 were to deviate and choose any cognition $\rho'_2 \geq \rho_1$, her stage-2 optimal action would be the same as if the deviation was observed by player 2. That is, for any $\omega'_2 \equiv (\omega^1, \dots, \omega^{\rho'_2})$, player 2's stage-2 optimal action following the deviation to ρ'_2 would coincide with the equilibrium action

$$a_2^{\rho_1, \rho'_2}(\omega'_2) = \left[1 + \sum_{k=1}^{\rho_1} \omega^k \right] / (1 - \beta) + \sum_{k=\rho_1+1}^{\rho'_2} \omega^k$$

that player 2 would select if the deviation was observable. Player 2's ex-ante gross expected payoff following the deviation to $\rho'_2 \geq \rho_1$ is then equal to

$$V_2(\rho'_2; \rho) = V_2(\rho'_2; (\rho_1, \rho'_2)).$$

and coincides with her equilibrium payoff under the profile (ρ_1, ρ'_2) . Likewise, when player 2 deviates to cognition $\rho'_2 < \rho_1$, her stage-2 optimal action

$$\frac{1 + \sum_{k=1}^{\rho'_2} \omega^k}{1 - \beta}$$

and her ex-ante gross expected payoff $V_2(\rho'_2; \rho) = V_2(\rho'_2; (\rho_1, \rho'_2))$ coincide with the equilibrium levels if the deviation was observable. Again, this follows from the fact that player 2 (a) expects both players to respond to all dimensions $(1, \dots, \rho'_2)$ as if they were commonly known, and (b) has no information to predict player 1's response to any dimension above ρ'_2 .

Equipped with the above observations, we now turn to expectation conformity in these games.

Proposition 4 (sparsity-EC). *(a) Consider any pair of cognitive profiles $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ and $\rho = (\rho_1, \rho_2)$ such that $\hat{\rho}_2 > \rho_2 \geq \hat{\rho}_1 > \rho_1$. Irrespective of whether $\beta > 0$ (strategic complementarity) or $\beta < 0$ (strategic substitutability), UEC holds for such profiles (strictly for player 1, weakly for player 2), whereas ID holds as an equality for both players. As a result, EC holds strictly for player 1, but only weakly for player 2. (b) Next, consider any pair of cognitive profiles $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ and $\rho = (\rho_1, \rho_2)$ such that $\hat{\rho}_2 = \hat{\rho}_1 > \rho_2 = \rho_1$. UEC holds as an equality for these profiles, whereas ID holds if $\beta > 0$ but does not hold if $\beta < 0$ (that is, $\beta \Gamma_i^{ID}(\rho, \hat{\rho}) > 0$, $i = 1, 2$).*

Consider first cognitive profiles $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ and $\rho = (\rho_1, \rho_2)$ such that $\hat{\rho}_2 > \rho_2 \geq \hat{\rho}_1 > \rho_1$. As explained above, for the player who is expected to engage in more cognition (player 2), the value of exploring more dimensions is invariant in whether the opponent (player 1) expects her to explore $\hat{\rho}_2 \geq \rho_1$ or $\rho_2 \geq \rho_1$ dimensions. This is because player 1's stage-2 equilibrium action is the same under both expectations. For player 1, instead, the value of exploring more dimensions is larger when the opponent expects her to explore more dimensions. This is because player 2 changes her response to the dimensions $(\rho_1 + 1, \dots, \hat{\rho}_1)$ as a function of whether or not she expects player 1 to also learn these dimensions. In particular, player 2 increases her responses to such dimensions when the stage-2 actions are complements, whereas she reduces her response to such dimensions when the stage-2 actions are strategic substitutes. In either case, player 1 benefits from the adjustment in player 2's response. Hence, no matter the sign of β , UEC holds for these profiles (strictly for player 1, weakly for player 2).

Next, consider ID. For player 1, the value of expanding her cognition is invariant in player 2's depth of knowledge. This is because, as explained above, her best estimate of player 2's information is her own information, no matter how far ahead player 2 is in the exploration of the state. Likewise, when $\hat{\rho}_2, \rho_2 \geq \hat{\rho}_1, \rho_1$, the gross value that player 2 assigns to exploring the dimensions $(\rho_2 + 1, \dots, \hat{\rho}_2)$ is invariant to whether player 1 selects cognition $\hat{\rho}_1$ or cognition ρ_1 , for, in both cases, player 2 expects to be the sole learner of the dimensions $(\rho_2 + 1, \dots, \hat{\rho}_2)$. Hence, ID holds (but only weakly, as an equality) for such profiles.

Next, consider profiles $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ and $\rho = (\rho_1, \rho_2)$ such that $\hat{\rho}_2 = \hat{\rho}_1 > \rho_2 = \rho_1$. In this case, the value that each player assigns to expanding her cognition is invariant to whether or not the opponent expects her to explore more dimensions: the opponent's action is invariant to such expectations. This is because, as explained above, the opponent reasons, and then acts, as if none of the dimensions $\kappa > \rho_j$ existed. Hence, UEC holds as an equality for these profiles. Fixing the opponent's expectations about her own cognition, we then have that the value a player assigns to going deeper in the exploration of the state is larger when the opponent also goes deeper under strategic complements but lower under strategic substitutes. This is because the value of learning more dimensions is higher under joint learning than under sole learning when actions are complements, whereas the opposite is true when actions are substitutes.

When paired with the results in Proposition (1), the results in part (a) in the above proposition thus suggest that these games may admit multiple asymmetric equilibria when cost functions are asymmetric across players, both in the case of complements and in the case of substitutes. The results in part (b), instead, suggest uniqueness of a symmetric equilibrium when the cost functions are symmetric across players in case the stage-2 actions are substitutes, and multiplicity of symmetric equilibria when the stage-2 actions are complements. The next two propositions verify that these conjectures are indeed correct and provide a more detailed account of the type of equilibrium multiplicity these games are prone to.

Proposition 5 (sparsity–complements). *Suppose that the stage-2 actions are strategic complements ($0 < \beta < 1$) and that cognitive costs are symmetric across players and take the linear form $C_i(\rho_i) \equiv \sum_{k=1}^{\rho_i} c_k$, with $c \equiv (c^k)_{k=1}^K \in \mathbb{R}_{++}^K$ such that σ_k^2/c_k is strictly decreasing in k . The following are true:*

- All (pure-strategy) equilibria are symmetric.
- Let

$$\underline{k} \equiv \min \left\{ k \mid \sigma_k^2 \leq c_k \right\} \quad \text{and} \quad \bar{k}(\beta) \equiv \max \left\{ k \mid \frac{\sigma_k^2}{(1-\beta)^2} \geq c_k \right\}.$$

Any level of cognition $k^ \in [\underline{k}, \bar{k}(\beta)]$ can be part of a symmetric (pure-strategy) equilibrium, and only these levels can be sustained in a symmetric (pure-strategy) equilibrium.*

- Suppose that there are no external payoff effects, meaning that individual payoffs u_i depend on other agents' actions a_{-i} only through the effect that the latter have on individual best responses.²² Then the (pure-strategy) equilibria are Pareto ranked, with the players' net payoff increasing in the equilibrium depth of knowledge $k^* \in [\underline{k}, \bar{k}(\beta)]$.

²²When payoffs are as in (7), this occurs when $\psi(a_{-i}, \omega) = h(\omega) + \beta \bar{a}_{-i}^2$, where $h(\omega)$ is an arbitrary function of ω . In this case, the externality $\beta \bar{a}_{-i}^2$ originating in the second term in (7) is perfectly offset by the corresponding term in

- *The equilibria described above are also equilibria when cognition is flexible, i.e., when players can explore a dimension k without having explored all dimensions $k' < k$.*

That, when stage-2 actions are complements and costs are symmetric across players, all equilibria are symmetric follows from the fact that the net benefit of learning the k -th dimension is larger for a player who knows the opponent also learns that dimension than for a player who expects to be the sole learner of that dimension. This property in turn implies that if an asymmetric equilibrium existed in which player 2 learns more dimensions than player 1, then player 1 would have a profitable deviation. Hence there do not exist equilibria with asymmetric levels of cognition. It is also easy to see that, when the two players select the same cognition, they then play the same stage-2 actions. Hence asymmetric equilibria where the asymmetry originates solely in the stage-2 actions do not exist either.

That, in any symmetric equilibrium, the cognitive level k^* must exceed \underline{k} follows from the fact that, if this was not the case, then a player would have incentives to deviate and learn the extra dimension $k^* + 1$ (the net value of being the sole learner of dimension k is equal to $\sigma_k^2 - c_k$, which is strictly positive for any $k < \underline{k}$).

Likewise, that the equilibrium depth of knowledge k^* must not exceed $\bar{k}(\beta)$ follows from the fact, if this condition was not satisfied, a player would have incentives to stop his exploration of the state at dimension $k^* - 1$ (the net benefit of learning the k^* -dimension when both players are expected to learn k^* dimensions is $\sigma_k^2/(1 - \beta)^2 - c_k$, which is negative for $k > \bar{k}(\beta)$). Furthermore, because σ_k^2/c_k is decreasing, the condition $k^* \in [\underline{k}, \bar{k}(\beta)]$ is not only necessary but also sufficient for a symmetric equilibrium to exist with depth of cognition equal to k^* (the result follows from the characterization of the payoff functions $V_i(\rho'_i; \rho)$ provided above).

That, in the absence of external payoff effects, equilibria are Pareto ranked follows from the fact that, in any symmetric equilibrium with depth of knowledge k^* , the equilibrium payoffs (net of terms that depend on ω but that are invariant in the players' behavior) are given by

$$V_i(k^*; (k^*, k^*)) - C(k^*) = \frac{1 + \sum_{k=1}^{k^*} \sigma_k^2}{(1 - \beta)^2} - \sum_{k=1}^{k^*} c_k,$$

which are increasing in k^* . A larger depth of knowledge permits the players to better align their actions with the underlying state while also increasing, on average, the return to the opponent's actions. Finally, that equilibria in the game in which cognition naturally proceeds in a sequential order (meaning that players have to learn the k -th dimension before they can learn the $k + 1$ -th one) are also equilibria in the game in which the players can select which dimensions to learn in any order of their choice follows from the fact that the benefit σ_k^2/c_k of learning dimension k , relative to its cost, is decreasing in k , which implies that the first dimensions are the most salient ones, both when a player is the sole learner, and when he expects her opponent to also learn such dimensions. Also note that the number of dimensions learnt in the maximally informative equilibrium $\bar{k}(\beta)$ is increasing in β .

The next result provides a similar characterization for the case of strategic substitutes.

ψ . As a result, individual payoffs u_i depend on other agents' actions a_{-i} only through the effect that the latter have on individual best responses.

Proposition 6 (sparsity–substitutes). *Suppose that the stage-2 actions are strategic substitutes ($-1 < \beta < 0$) and that cognitive costs are symmetric across players and take the linear form $C_i(\rho_i) \equiv \sum_{k=1}^{\rho_i} c_k$, with $c \equiv (c^k)_k \in \mathbb{R}_{++}^K$ such that σ_k^2/c_k is strictly decreasing. The following are true:*

- *A (pure-strategy) symmetric equilibrium exists if and only if there exists a cognitive level k^* satisfying*

$$\frac{\sigma_{k^*+1}^2}{c_{k^*+1}} \leq 1 \leq \frac{\sigma_{k^*}^2}{(1-\beta)^2 c_{k^*}}. \quad (9)$$

When a symmetric (pure-strategy) equilibrium exists, it is unique and its cognitive level k^ satisfies (9).*

- *There may exist asymmetric (pure-strategy) equilibria. In any such equilibrium, the two players' cognitive levels ρ_1 and ρ_2 belong to $[\underline{k}(\beta), \bar{k}]$ where*

$$\underline{k}(\beta) \equiv \min \left\{ k \mid \sigma_{k+1}^2 (1+\beta)^2 \leq c_{k+1} \right\} \quad \text{and} \quad \bar{k} \equiv \max \left\{ k \mid \sigma_k^2 \geq c_k \right\}$$

and are such that

$$\frac{\sigma_{\rho_1+1}^2 (1+\beta)^2}{c_{\rho_1+1}} \leq 1 \leq \frac{\sigma_{\rho_1}^2}{(1-\beta)^2 c_{\rho_1}} \quad (10)$$

and

$$\frac{\sigma_{\rho_2+1}^2}{c_{\rho_2+1}} \leq 1 \leq \frac{\sigma_{\rho_2}^2}{c_{\rho_2}}. \quad (11)$$

Player 1's equilibrium payoff is increasing in her own depth of knowledge ρ_1 and invariant in player 2's depth of knowledge ρ_2 . Player 2's equilibrium payoff is decreasing in player 1's depth of knowledge ρ_1 and, in case there are multiple solutions ρ_2 to the double inequality in (11), is invariant in ρ_2 .²³ The sum of the two players' equilibrium payoffs is maximal under the (pure-strategy) equilibrium featuring the lowest cognition for player 1 (the one learning the smallest number of dimensions).

That sparsity games with substitutes feature at most one symmetric equilibrium follows from the fact that the gross benefit of learning any of the dimensions is larger for a player who expects to be the sole learner of that dimension than for a player who expects the other player to also learn the same dimension. Observe that $\sigma_{k^*+1}^2 - c_{k^*+1}$ is the value of exploring dimension $k^* + 1$ when a player expects to be the sole learner of such dimension, whereas $\sigma_{k^*}^2 / (1-\beta)^2 - c_{k^*}$ is the value of learning dimension k^* when both players are expected to learn that dimension. The double inequality in Condition (9) then guarantees that, starting from a situation where both players learn the same number of dimensions, no player has a profitable deviation. Clearly, there is at most one dimension k^* satisfying the aforementioned double inequality. Hence, the game admits at most one symmetric (pure-strategy) equilibrium and, when such an equilibrium exists, the equilibrium cognition k^* is given by the unique solution to the double inequality in Condition (9).

²³This is because multiple solutions to the two inequalities in (11) are possible only if either $\frac{\sigma_{\rho_2+1}^2}{c_{\rho_2+1}} = 1$ or $\frac{\sigma_{\rho_2}^2}{c_{\rho_2}} = 1$. In the first case, both ρ_2 and $\rho_2 + 1$ can be sustained in equilibrium. In the second case, both ρ_2 and $\rho_2 - 1$ can be sustained in equilibrium. In either case, under the largest solution, player 2 is indifferent between learning all the equilibrium dimensions and stopping one dimension earlier.

Next, consider asymmetric (pure-strategy) equilibria. In any such equilibrium, the value to player 2 of increasing (alternatively, decreasing) her cognition by one dimension is invariant in player 1's depth of knowledge. Hence, the number of dimensions explored by player 2 in any such equilibrium must satisfy the double inequality in (11). For player 1, instead, the number of dimensions explored in equilibrium depends on what player 2 expects her to do. From the discussion preceding Proposition 4, observe that the term $\sigma_{\rho_1+1}^2(1+\beta)^2 - c_{\rho_1+1}$ in the left-hand side of (10) is the net benefit of exploring the (ρ_1+1) -th dimension when player 2 also explores such dimension but does not expect player 1 to explore it, whereas the term $\sigma_{\rho_1}^2/(1-\beta)^2 - c_{\rho_1}$ in the right-hand side of (11) is the net benefit of exploring the ρ_1 -th dimension when player 2 also explores such dimension and expects player 1 to do the same. The double inequality, along with the monotonicity of the benefits/cost ratio σ_k^2/c_k in k , then implies that player 1 does not have profitable deviations.

Note that the benefit $\sigma_{\rho_1+1}^2(1+\beta)^2 - c_{\rho_1+1}$ of exploring the (ρ_1+1) -th dimension when not expected to do so is smaller than the benefit $\sigma_{\rho_1+1}^2/(1-\beta)^2 - c_{\rho_1+1}$ or exploring the same dimension when player 2 expects player 1 to learn the (ρ_1+1) -th dimension. This is because, in both cases, player 2 responds to such dimension, but, in the former case, player 2 does not adjust her action to accommodate for player 1's enhanced knowledge of the state. Hence, for the player who is behind, the marginal value of expanding her knowledge depends on what the other player expects her to do, which is the origin of the multiplicity of asymmetric equilibria.²⁴

The ranking of the two players' payoffs across the various asymmetric (pure-strategy) equilibria directly reflects the fact that (a) the value of learning more dimensions when expected to do so is always positive (thus explaining why player 1's equilibrium payoff is increasing in her depth of knowledge ρ_1), and (b) the value of learning a dimension is always higher when a player is a sole learner of that dimension than when both players learn it (thus explaining why player 2's equilibrium payoff is decreasing in player 1's depth of knowledge).

Finally, that the sum of the two players' equilibrium payoffs

$$\frac{1 + \sum_{k=1}^{\rho_1} \sigma_k^2}{(1-\beta)^2} + \sum_{k=\rho_1+1}^{\rho_2} \sigma_k^2 - 2\sum_{k=1}^{\rho_1} c_k - \sum_{k=\rho_1+1}^{\rho_2} c_k$$

is maximal under the equilibrium featuring the lowest cognitive level for player 1 follows from the fact that the net benefit $\sigma_k^2/(1-\beta)^2 - c_k$ that player 1 derives from learning an extra dimension is always smaller than the loss $\sigma_k^2 - \sigma_k^2/(1-\beta)^2$ that player 2 incurs when player 1 learns the extra dimension.

4.2 Noisy information acquisition

We now turn to environments in which cognition takes the (perhaps more familiar) form of noisy information acquisition. We start by considering games where players receive additive signals about an exogenous payoff state, and then consider games where players "spy" on each other, that is, where their cognition determines the precision of the additive signals they receive about their opponents' primitive signals.

²⁴That the game admits multiple asymmetric equilibria is consistent with the results in Propositions 1 and 4.

4.2.1 Learning about payoff states

Payoffs continue to be given by (7). Because the multi-dimensionality of the state is not important when the agents receive a noisy signal of the state, to simplify the derivations, assume that $K = 1$ (the state ω is thus uni-dimensional) and then let $g(\omega) = \omega$. Also, disregard the externality term $\psi(a_{-i}, \omega)$ in (7) as it plays no role in the computation of the individual best responses.

To further simplify the analysis, then assume that ω is drawn from an improper uniform prior over the entire real line (as it will become clear in a moment, hierarchies of beliefs and expected payoffs are well defined, despite the impropriety of the prior). Correlated (exogenous) variation in the players' beliefs about ω is then captured by the two players observing a common signal

$$y = \omega + \varepsilon,$$

with ε drawn from a Normal distribution with mean 0 and variance h^{-1} .

The result of each player's cognitive activity is instead summarized in a private signal

$$x_i = \omega + \eta_i,$$

with η_i drawn from a Normal distribution with mean 0 and variance ρ_i^{-1} , with (η_1, η_2) drawn independently across the two players and independently from (ω, ε) .²⁵

Now let $s_i \equiv (x_i, y)$, $i = 1, 2$. In the Appendix, we show that when, in the stage-2 game, player i expects player j to follow a strategy that, for any $s_j = (x_j, y)$ selects with probability one an action

$$a_j(s_j) = m_j x_j + (1 - m_j) y$$

then, given her own cognitive choice ρ_i , player i 's best response consists in following a strategy that, for each $s_i = (x_i, y)$, also selects with probability one an action $a_i(s_i) = m_i x_i + (1 - m_i) y$ that is a convex linear combination of x_i and y , with

$$m_i = \frac{(1 - \beta)\rho_i}{\rho_i + h} + \frac{\beta\rho_i}{\rho_i + h} m_j.$$

Fixing the two players' cognitive postures $\rho = (\rho_1, \rho_2)$, we then show that there exists a unique linear continuation equilibrium and this equilibrium is such that, for each $s_i = (x_i, y)$, $a_i^\rho(s_i) = m_i^\rho x_i + (1 - m_i^\rho) y$, with

$$m_i^\rho = (1 - \beta)\rho_i \frac{\rho_j(1 + \beta) + h}{(\rho_i + h)(\rho_j + h) - \beta^2 \rho_i \rho_j}. \quad (12)$$

Now suppose that the two players are expected to select the cognitive profile $\rho = (\rho_i, \rho_j)$ and player i deviates and selects cognition ρ'_i (and then chooses her period-2 action optimally). As we show in the Appendix, player i 's ex-ante expected payoff, gross of any cognitive cost, is then equal to

$$V_i(\rho'_i; \rho) = -\frac{\left(m_i^{\rho'_i; \rho}\right)^2}{\rho'_i} - \beta \frac{\left(m_j^\rho\right)^2}{\rho_j} - \left[1 - \beta + \left(m_i^{\rho'_i; \rho}\right)^2 - 2(1 - \beta)m_i^{\rho'_i; \rho} + \beta \left(m_j^\rho\right)^2 - 2\beta m_i^{\rho'_i; \rho} m_j^\rho\right] \frac{1}{h},$$

²⁵As shown in Bergemann and Morris (2013), any joint (Gaussian) distribution between (a_1, a_2, ω) can be generated by a combination of a perfectly public and two perfectly private Gaussian additive signals.

where

$$m_i^{\rho'_i; \rho} = \frac{(1 - \beta)\rho'_i}{\rho'_i + h} + \frac{\beta\rho'_i}{\rho'_i + h} m_j^\rho \quad (13)$$

is the sensitivity of player i 's period-2 action to his private information x_i when the two players are expected to select the cognitive profile $\rho = (\rho_i, \rho_j)$ and, instead, player i selects cognition ρ'_i . We then have the following result (the proof is in the Appendix):

Proposition 7 (learning about exogenous states). *Let $\rho = (\rho_1, \rho_2)$ and $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ be two arbitrary cognitive profiles. UEC always holds for these profiles, irrespective of whether the stage-2 actions are strategic complements ($\beta > 0$) or strategic substitutes ($\beta < 0$). ID holds if and only if $\beta(\hat{\rho}_i - \rho_i)(\hat{\rho}_j - \rho_j) \geq 0$. Finally, expectation conformity holds if and only if*

$$\beta(\hat{\rho}_j - \rho_j)(\hat{\rho}_i - \rho_i)(\hat{\rho}_i + h)[\rho_i(1 + \beta) + h] + \beta^2(\hat{\rho}_i - \rho_i)^2 \hat{\rho}_j[\rho_j(1 + \beta) + h] \geq 0 .$$

The result in the first part of the proposition says that, holding fixed the precision of player j 's private signal, the value to player i of increasing the precision of her private signal is higher when player j expects her to acquire a more precise private signal. That is, the game always satisfies UEC. To see this, consider first the case where the stage-2 actions are strategic complements ($\beta > 0$). When player j expects player i to acquire a more precise signal, she also expects player i to select an action in the stage-2 game that is more sensitive to player i 's private signal. Because $\beta > 0$, player j 's stage-2 action is then more sensitive to player j 's own private signal (use (12) to verify that m_j^ρ is increasing in ρ_i) which in turn increases player i 's incentives to acquire a more precise private signal.

Next, consider the case where the stage-2 actions are strategic substitutes ($\beta < 0$). Again, when player j expects player i to acquire a more precise private signal, she also expects player i 's stage-2 action to be more sensitive to player i 's private signal. Because $\beta < 0$, player j 's best response is then to select a stage-2 action that is less sensitive to player j 's own private signal (again, use (12) to verify that, when $\beta < 0$, m_j^ρ is decreasing in ρ_i). Because actions are strategic substitutes, that player j 's stage-2 action responds less to her private signal in turn implies that player i 's incentives to acquire a more precise private signal are stronger. Hence, irrespective of whether $\beta > 0$ or $\beta < 0$, holding fixed player j 's cognition ρ_j , player i 's incentives to acquire more precise private information are always higher when player j expects her to acquire more precise private information.

The second part of the proposition says that, holding player j 's expectation about player i 's cognition fixed, the value to player i of acquiring more precise private information is higher when either (a) player j also acquires more precise private information and actions are strategic complements ($\beta > 0$), or (b) player j acquires less precise private information and actions are strategic substitutes. This is because, irrespective of whether $\beta > 0$ or $\beta < 0$, when player j acquires more precise private information, she then always responds more to it (use (12) to verify that m_j^ρ is always increasing in ρ_j , irrespective of the sign of β). Player i 's incentives to acquire more precise information are thus higher when either $\beta > 0$ and $\hat{\rho}_j \geq \rho_j$, or when $\beta < 0$ and $\hat{\rho}_j < \rho_j$ (they are lower when $\beta(\hat{\rho}_j - \rho_j) \leq 0$).

Finally, the last part of the proposition says that player i 's incentives to acquire more precise information (that is, to choose $\hat{\rho}_i > \rho_i$) are stronger under the profile $\hat{\rho} = (\hat{\rho}_i, \hat{\rho}_j)$ than under the profile $\rho = (\rho_i, \rho_j)$ when, other things equal, $\hat{\rho}_i - \rho_i$ is large relative to $\hat{\rho}_i \rho_i$. For example, when public information is imprecise, that is, when $h \rightarrow 0$, EC holds if and only if

$$\beta(\hat{\rho}_j - \rho_j)\hat{\rho}_i \rho_i + \beta^2(\hat{\rho}_i - \rho_i)\hat{\rho}_j \rho_j \geq 0.$$

The literature on noisy information acquisition in linear-quadratic-Gaussian games such as the one considered here has noticed that such games may or may not feature multiple equilibria, depending on the structure of the information cost C . For example, when $\beta > 0$, Hellwig and Veldkamp (2009) noticed that such games typically admit multiple (symmetric) equilibria when players have access to a finite set of information sources, whereas Myatt and Wallace (2012) have noticed that the same games typically admit a unique symmetric equilibrium when players have access to a continuum of information sources, with the cost of acquiring a more precise signal vanishing as the precision of the signal also vanishes. The results in Proposition 7 above contribute a different angle. First, they shed light about the players' incentives to conform to their opponents' expectations, for *arbitrary cognitive profiles* $\rho = (\rho_i, \rho_j)$ and $\hat{\rho} = (\hat{\rho}_i, \hat{\rho}_j)$, which are not restricted to be symmetric across the players. More importantly, they help identify the underlying forces contributing to equilibrium multiplicity and in particular whether the latter originates in the players benefiting from conforming to the other players' expectations about their own cognition (UEC) or in their expectations about the opponents' own cognitive postures (ID).

4.2.2 Learning about other players' beliefs: Espionage

Next, consider a situation in which the information the players collect permits them to learn not only about fundamental payoff-relevant variables (captured by ω) but also about non-fundamental variation in other players' beliefs. An example of such a situation is when players “spy” on each other, that is, acquire information about other players' primitive signals. We capture such a situation as follows. The payoffs continue to be given by (7), with $K = 1$, $g(\omega) = \omega$, and $\psi(\omega, a_{-i}) = 0$, and with ω drawn from an improper uniform distribution over the entire real line. Contrary to what assumed in the previous subsection, there is no exogenous public information y (the conclusions below extend to settings with exogenous public information, but because the latter plays no fundamental role in the analysis, we omit it to ease the derivations). Each player $i = 1, 2$ is endowed with an exogenous “primary” signal

$$s_i^P = \omega + \varepsilon_i$$

with ε_i drawn from a standard Normal distribution. Such a signal captures the player's primitive private information. In addition to s_i^P , player i can acquire a “secondary” signal

$$s_i^S = s_j^P + \eta_i = \omega + \varepsilon_j + \eta_i$$

where the noise η_i is drawn from a Normal distribution with mean 0 and precision ρ_i . This secondary signal is thus information that the player collects about the opponent's primary signal s_j^P . It contains information not only about ω but also about the noise ε_j in the opponent's primary signal. It can be interpreted as the result of industrial espionage, as in Kozlovskaya (2018) and in Adriani and Sonderegger (2020). More broadly, it can be thought of as the result of various activities that help the player interpret the opponent's view of the game (see also Angeletos and La'O (2013) for a model in which agents receive information about the noise in other agents' beliefs, but where such information is exogenous, and Calvo-Armengol et al. (2015) and Sethi and Yildiz (2016, 2018) for models in which players choose the information they receive from other players). The variables $(\omega, \varepsilon_1, \varepsilon_2, \eta_1, \eta_2)$ are mutually independent.

As in the previous subsection, suppose that, in the stage-2 game, player i expects player j to

follow a strategy that, for any $s_j \equiv (s_j^P, s_j^S)$, selects with probability one the action

$$a_j(s_j) = m_j s_j^P + (1 - m_j) s_j^S$$

Given ρ_i , for any $s_i \equiv (s_i^P, s_i^S)$, player i 's best response then consists in selecting with probability one the action $a_i(s_i) = m_i s_i^P + (1 - m_i) s_i^S$ with

$$m_i = \frac{1 + \rho_i(1 + \beta) - 2\beta\rho_i m_j}{1 + 2\rho_i}. \quad (14)$$

When the game features strategic complements ($\beta > 0$), the higher the sensitivity of player j 's period-2 action to her primary signal (that is, the larger m_j is), the smaller the weight that player i 's best response assigns to player i 's primary signal (equivalently, the larger the weight $1 - m_i$ to the secondary signal s_i^S). The opposite is true in case of strategic substitutes ($\beta < 0$).

Fixing the two players' cognitive postures $\rho = (\rho_1, \rho_2)$, we then have that there exists a unique stage-2 linear continuation equilibrium and is such that the equilibrium actions are given by $a_i^\rho(s_i) = m_i^\rho s_i^P + (1 - m_i^\rho) s_i^S$, $i = 1, 2$, with

$$m_i^\rho = \frac{1 + 2\rho_j + \rho_i(1 - \beta) + 2\rho_i\rho_j(1 - \beta^2)}{1 + 2(\rho_i + \rho_j) + 4\rho_i\rho_j(1 - \beta^2)} \quad (15)$$

Next, let

$$m_i^{\rho'_i; \rho} = \frac{1 + \rho'_i(1 + \beta) - 2\beta\rho'_i m_j^\rho}{1 + 2\rho'_i}$$

denote the sensitivity of player i 's stage-2 strategy to her primary signal when the two players are expected to choose cognition $\rho = (\rho_i, \rho_j)$ and, instead, player i chooses cognition ρ'_i . Using the fact that

$$\begin{aligned} V_i(\rho'_i; \rho) &= -(1 - \beta) \left(m_i^{\rho'_i; \rho} \right)^2 - (1 - \beta) \left(1 - m_i^{\rho'_i; \rho} \right)^2 \left(1 + \frac{1}{\rho'_i} \right) \\ &\quad - \beta \left[m_i^{\rho'_i; \rho} - \left(1 - m_j^\rho \right) \right]^2 - \beta \left[\left(1 - m_i^{\rho'_i; \rho} \right) - m_j^\rho \right]^2 - \beta \left(1 - m_i^{\rho'_i; \rho} \right)^2 \frac{1}{\rho'_i} - \beta \left(1 - m_j^\rho \right)^2 \frac{1}{\rho_j} \end{aligned}$$

we then have the following result (the proof is in the Appendix):

Proposition 8 (espionage). *Let $\rho = (\rho_i, \rho_j)$ and $\hat{\rho} = (\hat{\rho}_i, \hat{\rho}_j)$ be two arbitrary cognitive profiles. UEC always holds for these profiles, irrespective of whether the stage-2 actions are strategic complements ($\beta > 0$) or strategic substitutes ($\beta < 0$). ID holds if and only if $\beta(\hat{\rho}_i - \rho_i)(\hat{\rho}_j - \rho_j) \leq 0$.*

The first result (UEC) parallels the corresponding result in Proposition (7). Holding player j 's cognition fixed, the value to player i of spying on player j (more broadly, of investing in cognitive activities that increase her ability to understand player j 's primitive view of the game) is higher, the more player j expects player i to invest in cognition (that is, the more player j expects player i to spy on her). This is because, independently of the sign of β , when player i spies more, she then relies more on her secondary signal (use the formula in (15) to verify that m_i^ρ is decreasing in ρ_i , which means that the sensitivity $1 - m_i$ of player i 's stage-2 action to her secondary signal s_i^S is increasing in ρ_i). When $\beta > 0$, this induces player j to rely more on her primary signal—player j likes being

spied (use (14) applied to player j to verify that m_j is decreasing in m_i and hence increasing in $1 - m_i$). This makes it even more valuable for player i to rely on her secondary signal and hence to spy more. When, instead, $\beta < 0$, player j responds less to her primary signal when she expects player i to spy more. This is because player j wants to distance herself from player i when actions are strategic substitutes. That player j relies less on her primary signal, however, further boosts player i 's incentives to spy on player j , as player i can now learn player j 's information without ending up aligning her action much with player j 's. Thus UEC holds even when actions are substitutes.

Next, consider ID. In this case, the result is the opposite of the one in the previous proposition. When $\beta > 0$, the value to player i of spying on player j is higher, the less player j spies, whereas the opposite is true when $\beta < 0$. Recall that, no matter the sign of β , when player j spies more, she then relies more on her secondary signal and less on her primary one. When $\beta > 0$, this reduces player i 's incentives to spy on player j , whereas the opposite is true when $\beta < 0$. One can also identify conditions under which, when ID does not hold, UEC is nonetheless strong enough to guarantee EC, as well as conditions under which ID prevails over UEC, thus inducing a negative form of expectation conformity. The conditions, however, are less illuminative than those in Proposition (7) and hence are not further discussed here.

5 Manipulative Cognition

In many environments of interest, cognition has a *manipulative* dimension: a player's cognition impacts her opponents' understanding of the game. We start by considering situations in which players choose "frames," or other manipulative devices, to influence other players' recollection of information. These games belong in the broader class of strategic situations in which players engage in signal-jamming. We then consider games in which players invest to help other players' understand their own view of the game, as in the literature on noisy communication,²⁶ or to prevent other players from spying on them, as in the case of counter-espionage.

5.1 Framing and defensive memory management

Consider a situation in which a player (the persuader) tries to induce another player (the receiver) to act favorably to her, by manipulating the receiver's recollection of information relevant for a decision. The manipulation is done by means of "frames," that is, through the design of a contextual purchasing experience—see Salant and Siegel (2018) for various examples along these lines.

We capture such situations as follows. Player 2 (the receiver) has a payoff equal to

$$u_2(a_1, a_2, \omega) = -(a_2 - \omega)^2$$

where $a_2 \in \mathbb{R}$ is player 2's action and where $\omega \in \mathbb{R}$ is the underlying state of Nature. Player 1 (the persuader), instead, has a payoff

$$u_1(a_1, a_2, \omega) = a_2$$

²⁶See Dewatripont and Tirole (2005) and the references therein.

that is invariant in ω and in her own action, and increasing in player 2's action.²⁷ Hence, player 2 wants to “do the right thing” (i.e., align her action to the underlying state ω), whereas player 1 wants player 2 to take as high an action as possible (e.g., to increase her purchases of player 1's product, irrespectively of whether or not this is good for player 2). This structure has received considerable attention in the recent persuasion and information-design literature. Contrary to what typically assumed in this literature, though, here player 2 cannot commit to her choice of a frame (i.e., to her information structure).

Player 2, the receiver, is originally endowed with a primary (exogenous) signal $s_2^P = \omega + \varepsilon$ but recalls such a signal only imperfectly. Such a primary signal may represent the information a buyer received about a seller's product from exogenous sources, or past experiences. In such a context, a “frame” by player 1 is a device influencing player 2's ability to recollect her primary signal. Importantly, such a frame may operate asymmetrically across states, facilitating the recollection of information favorable to 1 relative to the less favorable one. The choice of a frame may also depend on the information that player 1 herself has about the state. However, because this channel is not essential to the results, we do not consider it here. Instead, we allow the receiver, player 2, to exert effort to increase her recollection of the primary information, thus reducing the effect of player 1's frame on her decision. We interpret such efforts broadly as “defensive memory management.” Allowing for such efforts also permits us to investigate whether ID holds in this context.

Let $\rho_1, \rho_2 \in \mathbb{R}_+$ and denote by $r(s_2^P; \rho)$ the probability that player 2 recalls her primary signal s_2^P when the two players engage in cognition $\rho = (\rho_1, \rho_2)$. Let $s_2^R \in \mathbb{R} \cup \{\emptyset\}$ denote player 2's recalled signal, with $s_2^R = \emptyset$ in case player 2 does not recall, and $s_2^R = s_2^P$ in case she does recall. Without loss of generality, then let $s_2^P = \omega$, with ω drawn from \mathbb{R} according to some cdf F .

For simplicity, assume that player 1's cognition (equivalently, her choice of a frame) increases uniformly the probability that player 2 recalls any positive signal and leaves it unaltered the probability that player 2 recalls any negative signal. Such a stark structure is not essential to the results. What matters is that the likelihood that the receiver recollects information that is more favorable to the persuader relative to the less favorable one is non-decreasing in the usage of frames. Formally, there exist non-negative and non-decreasing functions r^+ and r^- such that

$$r(\omega; \rho) = \begin{cases} r^-(\rho_2) & \text{if } \omega < 0 \\ r^+(\rho_1, \rho_2) & \text{if } \omega \geq 0. \end{cases}$$

Given the cognitive profile $\rho = (\rho_1, \rho_2)$, then let

$$\mathbb{E} [\tilde{\omega} | s_2^R; \rho] = \begin{cases} \omega & \text{if } s_2^R = \omega \\ \bar{\omega}(\rho) & \text{if } s_2^R = \emptyset \end{cases}$$

denote player 2's posterior expectation of the state (equivalently, of her optimal action) given the recalled information. Let $\omega^- = \mathbb{E} [\tilde{\omega} | \tilde{\omega} < 0]$ and $\omega^+ = \mathbb{E} [\tilde{\omega} | \tilde{\omega} \geq 0]$, where both expectations are under the prior distribution F . We then have that, in the absence of any recollection of her primitive information, player 2's expected value of ω is equal to

$$\bar{\omega}(\rho) = \frac{(1 - r^-(\rho_2))F(0)\omega^- + (1 - r^+(\rho_1, \rho_2))(1 - F(0))\omega^+}{(1 - r^-(\rho_2))F(0) + (1 - r^+(\rho_1, \rho_2))(1 - F(0))}.$$

²⁷The above payoff specification is thus essentially the same as the one in (7) in the previous section, with $\beta = 0$, $\psi(a_{-i}, \omega) = a_{-i}$, $g(\omega) = \omega$, and $|A_1| = 1$.

Note that $\bar{\omega}(\rho)$ is weakly decreasing in ρ_1 , that is, in the beliefs player 2 has about player 1's use of manipulative frames. It may be either increasing or decreasing in player 2's own cognition, ρ_2 . In particular, $\bar{\omega}(\rho)$ is decreasing in ρ_2 if $\frac{dr^-(\rho_2)}{d\rho_2} = \frac{\partial r^+(\rho_1, \rho_2)}{\partial \rho_2}$ and $r^+(\rho_1, \rho_2) \geq r^-(\rho_2)$, that is, if more cognition by player 2 has an equal effect on her ability to recollect positive and negative information, and if the likelihood that she recollects positive information is no smaller than the likelihood that she recollects negative information. On the other hand, $\bar{\omega}(\rho)$ is increasing in ρ_2 , when $\frac{dr^-(\rho_2)}{d\rho_2} > \frac{\partial r^+(\rho_1, \rho_2)}{\partial \rho_2}$ and $r^+(\rho_1, \rho_2) \cong r^-(\rho_2)$.

Given the quadratic loss function, for any cognitive profile ρ , and any recalled memory s_2^R , player 2's optimal action is equal to

$$a_2^\rho(s_2^R) = \mathbb{E} [\tilde{\omega} | s_2^R; \rho]$$

implying that, for any cognitive profile $\rho = (\rho_1, \rho_2)$ and any actual choice of frame ρ'_1 , player 1's ex-ante expected gross payoff when the two players are expected to engage in cognition $\rho = (\rho_1, \rho_2)$ and, instead, player 1 chooses ρ'_1 is equal to

$$V_1(\rho'_1; \rho) = F(0) [(1 - r^-(\rho_2))\bar{\omega}(\rho) + r^-(\rho_2)\omega^-] + (1 - F(0)) [(1 - r^+(\rho'_1, \rho_2))\bar{\omega}(\rho) + r^+(\rho'_1, \rho_2)\omega^+].$$

Similarly, given any cognitive profile $\rho = (\rho_1, \rho_2)$ and any actual choice ρ'_2 of memory management, player 2's ex-ante expected gross payoff when the two players are expected to engage in cognition $\rho = (\rho_1, \rho_2)$ and, instead, player 2 selects cognition ρ'_2 , is equal to

$$\begin{aligned} V_2(\rho'_2; \rho) &= -F(0)(1 - r^-(\rho'_2))\mathbb{E} [(\bar{\omega}(\rho_1, \rho'_2) - \tilde{\omega})^2 | \tilde{\omega} < 0] \\ &\quad - (1 - F(0))(1 - r^+(\rho_1, \rho'_2))\mathbb{E} [(\bar{\omega}(\rho_1, \rho'_2) - \tilde{\omega})^2 | \tilde{\omega} \geq 0]. \end{aligned}$$

The next result illustrates how cognitive expectations shape the players' incentives to engage into manipulative framing and to invest in defensive memory management in such environments.

Proposition 9 (framing and memory management). *Consider any pair of cognitive profiles $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ and $\rho = (\rho_1, \rho_2)$. UEC always holds for such profiles (weakly for the receiver, player 2, and strongly for the persuader, player 1). ID holds for player 1 (the persuader) if and only if*

$$[r^+(\hat{\rho}_1, \hat{\rho}_2) - r^+(\rho_1, \hat{\rho}_2)] [\omega^+ - \bar{\omega}(\hat{\rho}_1, \hat{\rho}_2)] \geq [r^+(\hat{\rho}_1, \rho_2) - r^+(\rho_1, \rho_2)] [\omega^+ - \bar{\omega}(\hat{\rho}_1, \rho_2)] \quad (16)$$

which is the case for example when (a) $\hat{\rho}_1 \geq \rho_1$, $\hat{\rho}_2 \geq \rho_2$, (b) r^+ is weakly supermodular, and (c) $\bar{\omega}$ is weakly decreasing in ρ_2 .

The persuader's incentives to engage into manipulative framing are stronger when she is expected to invest more into manipulative framing. This is because, the more the receiver expects the persuader to engage in manipulative framing, the more she interprets the lack of recollection of her primitive information as a signal of the state being unfavorable to the persuader. But then the stronger the incentives for the persuader to engage into manipulative framing to reduce the risk that the receiver does not recall.

Next, consider the receiver, player 2. Her optimal action depends only on her beliefs about player 1's manipulation and not on her belief about player 1's expectation of her own defensive cognition.

As a result, UEC also holds for player 2 but in the trivial sense of player 2’s incentives being invariant in player 1’s expectations about player 2’s cognition.

The second part of the proposition identifies a condition under which player 1’s payoff from engaging in manipulative framing are stronger when player 1 is expected to invest more in defensive memory management. The condition holds, for example, when the more player 2 invests in recollecting her primitive information, the larger the marginal effect of player 1’s manipulation on player 2’s recollection of positive information and, in the absence of any recollection, the lower player 2’s optimal action. Increasing differences for player 1 also holds when r^+ is submodular (that is, when the more player 2 invests in defensive cognition, the smaller the marginal effect of player 1’s manipulation on player 2’s recollection of positive information) provided that, in the absence of any recollection, player 2’s optimal action is significantly smaller when player 2 invests more in memory management than when she invests less.

Whether increasing differences holds for player 2 (the receiver) is more convoluted and depends on a complicated condition which we do not discuss here.

5.2 Other instances of signal jamming

The manipulative frames considered in the previous subsection are instances of “signal jamming,” akin to those studied in the industrial organization literature. For example, signal jamming occurs when a firm secretly cuts its price so as to reduce its rivals’ profits and induce them to believe that demand is low (or that costs are high) and exit the market. Cognitive traps are common in such games. In this subsection, we discuss a few other examples of signal jamming considered in the literature, without however getting much into the details.

(a) *Generalized career concerns.* In Holmström (1999)’s celebrated model, a worker exerts effort to convince a competitive labor market that her talent is high. The worker’s performance depends on her talent, which is unknown to the worker, her effort, and noise. When talent and effort are complements, such signal jamming often generates expectation conformity and expectation traps (see Dewatripont et al. 1999).²⁸ To see this, consider a generalized version of the career-concerns model in which both the worker and the competitive labor market can invest to influence the information the labor market receives about the worker’s talent. Let player 1 be the worker and player 2 the competitive labor market, and assume that payoffs are as in the previous subsection. Denote player 1’s effort by ρ_1 and player 2’s effort by ρ_2 . Given $\rho = (\rho_1, \rho_2)$, player 2 receives a signal

$$s_2 = A(\rho) + M(\rho)\omega + R(\rho)\varepsilon_2$$

about player 1’s talent ω , where A is an “additive” term akin to the one in Holmström (1999)’s original model, M is a “multiplicative” term akin to the one in Dewatripont et al. (1999), and R is a term capturing player 2’s ability to “recall,” as in the framing application in the previous subsection. Each of these functions are non-negative, with A and M non-decreasing, and R non-increasing. Holmström (1999)’s original model corresponds to $A(\rho) = \rho_1$ and $M(\rho) = R(\rho) = 1$, all $\rho = (\rho_1, \rho_2)$ (only the worker invests and effort has an additive effect on performance), whereas the multiplicative model of Dewatripont et al. (1999) corresponds to $M(\rho) = \rho_1$, $A(\rho) = 0$, and

²⁸See also Horner and Lambert (2019) for a more recent analysis of these games.

$R(\rho) = 1$, all $\rho = (\rho_1, \rho_2)$ (again, only the worker invests, but the impact of effort on performance now depends on the state ω). In either model, recall is exogenous.

Suppose that ω is normally distributed with mean $\omega_0 > 0$ and variance $1/h_\omega$, and that ε_2 is normally distributed with mean 0 and variance $1/h_\varepsilon$. Fixing $\rho = (\rho_1, \rho_2)$, for any s_2 , player 2's optimal action is then given by

$$a_2^\rho(s_2) = \mathbb{E}[\tilde{\omega}|s_2; \rho] = \left[\frac{M^2(\rho)h_\varepsilon}{M^2(\rho)h_\varepsilon + R^2(\rho)h_\omega} \right] \left[\frac{s_2 - A(\rho)}{M(\rho)} \right] + \left[\frac{R^2(\rho)h_\omega}{M^2(\rho)h_\varepsilon + R^2(\rho)h_\omega} \right] \omega_0.$$

Given $\rho = (\rho_1, \rho_2)$, for any actual choice ρ'_1 by player 1, player 1's ex-ante expected payoff (gross of the cognitive cost but net of all terms that do not depend on her actual choice ρ'_1) is equal to

$$V_1(\rho'_1; \rho) = \frac{M(\rho)h_\varepsilon}{M^2(\rho)h_\varepsilon + R^2(\rho)h_\omega} [M(\rho'_1, \rho_2)\omega_0 + A(\rho'_1, \rho_2)].$$

Likewise, given $\rho = (\rho_1, \rho_2)$, for any actual choice ρ'_2 by player 2, player 2's ex-ante expected payoff (gross of the cognitive cost but net of all terms that do not depend on her actual choice ρ'_2) is equal to

$$V_2(\rho'_2; \rho) = -\frac{M^2(\rho_1, \rho'_2)R^2(\rho_1, \rho'_2)h_\varepsilon + R^4(\rho_1, \rho'_2)h_\omega}{(M^2(\rho_1, \rho'_2)h_\varepsilon + R^2(\rho_1, \rho'_2)h_\omega)^2}$$

It is then easy to see that, when only player 1 invests, UEC (and hence EC) never obtains in Holmström (1999) additive model, where the optimal level of ρ'_1 is implicitly given by

$$C'_1(\rho'_1) = \frac{h_\varepsilon}{h_\varepsilon + h_\omega}$$

and is independent of the level ρ_1 expected by player 2. Instead, UEC (and hence EC) can easily obtain in the multiplicative model of Dewatripont et al. (1999) where the optimal level of ρ'_1 solves

$$C'_1(\rho'_1) = \frac{\rho_1 h_\varepsilon \omega_0}{\rho_1^2 h_\varepsilon + h_\omega}$$

and is increasing in ρ_1 for $\rho_1 \leq \sqrt{h_\omega/h_\varepsilon}$. Consistently with the results in Proposition 1, the equilibrium is thus unique in Holmström (1999), whereas multiple equilibria are possible in the multiplicative model of Dewatripont et al. (1999). Furthermore, when this is the case, the worker is better off in the low-effort equilibrium. This multiplicative version of this game is thus prone to expectation traps. Whether or not the above conclusions are robust to the possibility that player 2 (the labor market) also invests to influence the quality of the information received and/or to the possibility of endogenous recall depends on the specific assumptions one makes on the A , M , and R functions, as indicated in the proposition below. Let

$$L(\rho) \equiv M(\rho)\omega_0 + A(\rho),$$

$$G(\rho) \equiv \frac{M(\rho)h_\varepsilon}{M^2(\rho)h_\varepsilon + R^2(\rho)h_\omega},$$

and

$$Z(\rho) \equiv \frac{M^2(\rho)R^2(\rho)h_\varepsilon + R^4(\rho)h_\omega}{(M^2(\rho)h_\varepsilon + R^2(\rho)h_\omega)^2}$$

Proposition 10 (generalized-career-concerns). *Take any pair of profiles $\rho = (\rho_1, \rho_2)$ and $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ with $\hat{\rho}_i \geq \rho_i$, $i = 1, 2$. UEC trivially holds as equality for player 2. It holds for player 1 if the function G is non-decreasing in ρ_1 .²⁹ ID holds for player 1 if the function G is increasing in ρ_2 and the function L is supermodular. ID holds for player 2 if the function Z is submodular.*

(b) *Disclosure of semi-hard information.* Consider a trading game where the seller (player 1) with strictly positive probability knows the buyer’s value for her good, ω . For simplicity, suppose that ω can take only two values ω_1 and ω_2 , with $\omega_1 < \omega_2$, and that the seller does not value the good. By exerting effort, the seller can increase the probability that, when the state is ω_2 , the buyer understands that her value for the good is high. It is easy to verify that cognitive traps also arise naturally in this model. If the seller is expected by the buyer to exert substantial effort, the price p in the absence of persuasion (that is, when the buyer does not receive evidence that her value is ω_2) is low, in which case it is particularly profitable for the seller to invest to convince the buyer that the state is ω_2 when this is indeed the case. In case of multiplicity, the seller is better off in a lower-effort equilibrium. So expectation traps naturally arise in this context as well.

(c) *Intra-personal memory management.* Another class of signal-jamming games giving rise to expectation conformity and expectational traps is intra-personal memory-management games. This class of games describes situations in which a player receives information that she may try to remember or repress, to influence the behavior of her future selves. The individual may find herself in a trap, in which both repression and cognitive discipline are dictated by the expectations of her future selves, with distinct welfare implications.³⁰

5.3 Noisy information sharing and counter-espionage

In the persuasion environments of the last two subsections, only one player, player 2 acts in the stage-2 game, with the other player, player 1, acting as a persuader. We now consider a more symmetric environment in which both players have a non-degenerate action in the stage-2 game and can influence the opponent’s beliefs through their cognitive choices. The class of games we consider in this subsection is inspired by the literature on noisy communication (see, e.g., Dewatripont and Tirole, 2005, and Calvo-Armengol et al. 2015).

The stage-2 game has the same linear-quadratic structure as in Subsection 4.2, with payoffs as in (7), with $K = 1$, $g(\omega) = \omega$, $\psi(\omega, a_{-i}) = 0$, and with ω drawn from an improper uniform prior over the entire real line. As in Subsection (4.2.2), each player is endowed with an exogenous primary signal equal to

$$s_i^P = \omega + \varepsilon_i$$

²⁹Precisely, $\Gamma_1^{UEC}(\rho, \hat{\rho}) = 0$ if the function L is constant in ρ_1 . When, instead, L is strictly increasing in ρ_1 , then $\Gamma_1^{UEC}(\rho, \hat{\rho}) \geq 0$ if the function G is non-decreasing in ρ_1 (with the inequality strict when G is strictly increasing in ρ_1). When, instead, G is strictly decreasing in ρ_1 , then $\Gamma_1^{UEC}(\rho, \hat{\rho}) < 0$.

³⁰These games were introduced in Bénabou and Tirole (2002). See also Gottlieb (2014a,b). Dessi (2008) applies similar ideas in the context of cultural transmission with multiple agents. Bénabou (2013) and Bénabou and Tirole (2006) show how memory management and collective decisions interact to produce collective delusions.

and receives an endogenous secondary signal equal to

$$s_i^S = s_j^P + \gamma_j + \phi_i$$

However, contrary to what assumed in Subsection (4.2.2), a player's cognition now affects both the precision of the signal received from the opponent and the precision of the opponent's secondary signal. Specifically, ε_i is drawn from a standard Normal distribution, γ_j is drawn from a Normal distribution with mean zero and precision t_j , and ϕ_i is drawn from a Normal distribution with mean zero and precision l_i . As usual, the variables $(\omega, \varepsilon_1, \varepsilon_2, \gamma_1, \gamma_2, \phi_1, \phi_2)$ are mutually independent. Player i 's cognition $\rho_i = (l_i, t_i)$ now influences both player i 's ability to learn about player j 's information (through the l_i component) and player j 's ability to learn about player i 's information (through the t_i component). The l_i dimension thus controls for the self-directed component of player i 's cognition, as in the espionage game in Subsection (4.2.2). The novel t_i dimension, instead, controls for the manipulative component of i 's cognition. A higher t_i corresponds to a larger investment by player i in cognitive activities facilitating agent j 's ability to interpret agent i 's primitive information. Alternatively, l_i may represent player i 's investment in espionage whereas $1/t_i$ as her investment in "counter-intelligence," that is, in activities that make it more difficult for player j to "spy" on player i .

We start by illustrating how to accommodate for interactions between the two forms of cognition (self-directed and manipulative) and then specialize the analysis to the case where cognition is purely manipulative, so as to isolate the novel effects.

While it is natural to assume that player i 's cognitive cost is increasing in the effort l_i that she puts to interpret her opponent's information (alternatively, in the intensity of her spying activity), the dependence of C_i on t_i is context-specific. For example, if t_i is a proxy for the effort that player i puts in making herself understood, as in Dewatripont and Tirole (2005), then a higher t_i is likely to come with a higher cost for player i (explaining oneself neatly is costly). If, instead, t_i is the result of investments in counter-intelligence aimed at preventing the opponent from spying on player i 's information, then a higher t_i , by reflecting a smaller investment in such defensive measures, naturally comes with a lower cost.

Given $\rho_i = (l_i, t_i)$ and $\rho_j = (l_j, t_j)$, let

$$r_i \equiv \frac{t_j l_i}{t_j + l_i}$$

denote the endogenous precision of the total noise $\eta_i \equiv \phi_i + \gamma_j$ in player i 's secondary signal. The manipulative novel dimension is captured here by the fact that r_i depends also on t_j .

Following steps similar to those in Subsection (4.2.2), one can verify that, when, in the stage-2 game, player i expects player j to select, for each combination of primary and secondary signals $s_j = (s_j^P, s_j^S)$, an action $a_j(s_j) = m_j s_j^P + (1 - m_j) s_j^S$ that is a convex linear combination of s_j^P and s_j^S , her best response is to follow a stage-2 strategy that, for each $s_i = (s_i^P, s_i^S)$, selects with probability one an action $a_i(s_i) = m_i s_i^P + (1 - m_i) s_i^S$ that is also a convex linear combination of s_i^P

and s_i^S , with³¹

$$m_i = \frac{1 + r_i(1 + \beta) - 2\beta r_i m_j}{1 + 2r_i}.$$

Furthermore, given any pair of cognitive profiles $\rho = (\rho_1, \rho_2)$, with $\rho_i = (l_i, t_i)$, $i = 1, 2$, there exists a unique linear continuation equilibrium for the stage-2 game and is such that each player, for each $s_i = (s_i^P, s_i^S)$, selects with certainty the equilibrium action $a_i^\rho(s_i) = m_i^\rho s_j^P + (1 - m_i^\rho) s_j^S$ with

$$m_i^\rho = \frac{1 + 2r_j^\rho + r_i^\rho(1 - \beta) + 2r_i^\rho r_j^\rho(1 - \beta^2)}{1 + 2(r_i^\rho + r_j^\rho) + 4r_i^\rho r_j^\rho(1 - \beta^2)} \quad (17)$$

where

$$r_j^\rho \equiv \frac{t_i l_j}{t_i + l_j} \quad \text{and} \quad r_i^\rho \equiv \frac{t_j l_i}{t_j + l_i}.$$

One can also verify that, when the two players are expected to select the cognitive profiles $\rho = (\rho_i, \rho_j)$ and, instead, player i selects cognition $\rho'_i = (l'_i, t'_i)$, player i 's ex-ante expected gross payoff is then equal to

$$\begin{aligned} V_i(\rho'_i; \rho) = & -(1 + \beta) - 2 \left(m_i^{\rho'_i; \rho} \right)^2 - 2\beta \left(m_j^\rho \right)^2 + 2(1 + \beta) m_i^{\rho'_i; \rho} + 4\beta m_j^\rho (1 - m_i^{\rho'_i; \rho}) \\ & - \left(1 - m_i^{\rho'_i; \rho} \right)^2 \frac{1}{r_i^{\rho'_i; \rho_j}} - \beta \left(1 - m_j^\rho \right)^2 \frac{1}{r_j^{\rho'_i; \rho_j}} \end{aligned} \quad (18)$$

where

$$m_i^{\rho'_i; \rho} = \frac{1 + r_i^{\rho'_i; \rho_j}(1 + \beta) - 2\beta r_i^{\rho'_i; \rho_j} m_j^\rho}{1 + 2r_i^{\rho'_i; \rho_j}}$$

is the sensitivity of player i 's stage-2 action to his primary signal following the deviation to $\rho'_i = (l'_i, t'_i)$. Here

$$r_i^{\rho'_i; \rho_j} \equiv \frac{t_j l'_i}{t_j + l'_i} \quad \text{and} \quad r_j^{\rho'_i; \rho_j} \equiv \frac{t'_i l_j}{t'_i + l_j}$$

denote the precisions of the total noise $\eta_i \equiv \phi_i + \gamma_j$ and $\eta_j \equiv \phi_j + \gamma_i$ in the two agents' secondary signals that obtain when player j conforms to cognition $\rho_j = (l_j, t_j)$ whereas, instead, player i deviates to cognition $\rho'_i = (l'_i, t'_i)$.

As in the previous two subsections, the manipulative component of a player's cognition has important implications for expectation conformity and for the possibility of multiple equilibria and cognitive traps. To isolate its novel effects, but also to simplify some of the derivations, hereafter we shut down the self-directed component and focus on a situation in which cognition is purely manipulative.

Thus suppose that l_1 and l_2 are exogenously fixed and, without loss of generality, let l_1 and l_2 diverge to infinity (meaning that the noise terms ϕ_1 and ϕ_2 in the agents' secondary signals are identically equal to zero). The endogenous secondary signal $s_i^S = s_j^P + \gamma_j$ that each player receives is thus equal to the opponent's primary signal, s_j^P , augmented by a noise term γ_j entirely controlled by

³¹Note that the structure of these best responses is the same as in Subsection (4.2.2), except for the fact that the precision r_i of i 's secondary signal is now affected also by j 's cognition, which is the source of the manipulation under consideration here.

the opponent's cognitive effort. This situation is thus symmetrically opposite to the one in Subsection (4.2.2) where the noise in each player's secondary signal is controlled entirely by the player's own cognitive cost. While the analysis in Subsection (4.2.2) is meant to isolate the effects of espionage, the one in the present subsection is meant to isolate the effects of *counter-espionage*.

For any $\rho'_i = (t'_i, l'_i)$ and any $\rho = (\rho_i, \rho_j)$, with $\rho_i = (t_i, l_i)$ and $\rho_j = (t_j, l_j)$, we then have that

$$r_i^{\rho'_i, \rho_j} = t_j \quad \text{and} \quad r_j^{\rho'_i, \rho_j} = t'_i.$$

Hence, we can simplify the notation and interpret each player's cognition directly as the choice of the precision of the noise γ_j in the opponent's secondary signal. That is, we can let $\rho_i = r_j$, $i, j = 1, 2$, $j \neq i$. As mentioned above, we can then interpret a smaller ρ_i as a higher investment in counter-espionage, that is, in activities that make it difficult for the opponent to spy on a player's primitive information.

Because cognition is covert (that is, player j cannot observe deviations in cognition by player i) and because any deviation by player i has no direct effect on player i ' information, we have that, for any $\rho = (\rho_i, \rho_j)$, the sensitivity

$$m_i^{\rho'_i, \rho} = \frac{1 + \rho_j(1 + \beta) - 2\beta\rho_j m_j^\rho}{1 + 2\rho_j}$$

of player i 's action to her primary signal in the stage-2 game, following a deviation to cognition ρ'_i , is invariant in ρ'_i . We then have the following result (the proof is in the Appendix):

Proposition 11 (counter-espionage). *Let $\rho = (\rho_i, \rho_j)$ and $\hat{\rho} = (\hat{\rho}_i, \hat{\rho}_j)$. UEC holds for such cognitive profiles if $\beta > 0$, and does not hold if $\beta < 0$. Irrespective of the sign of β , ID holds if and only if $(\hat{\rho}_i - \rho_i)(\hat{\rho}_j - \rho_j) \leq 0$.*

Contrary to the spying game considered in Subsection (4.2.2), whether or not UEC holds in the counter-espionage game under consideration depends on whether the stage-2 actions are strategic complements or strategic substitutes. To gather some intuition, fix the precision ρ_j of the noise γ_j in the secondary signal $s_i^S = s_j^P + \gamma_j$ that player i expects to receive from player j . When player j expects player i to pass on a more precise signal (that is, when $\hat{\rho}_i \geq \rho_i$), in the stage-2 game, irrespective of the sign of β , player j then responds more to her secondary signal and less to her primary signal. To see this, use (17) to note that, given $\rho = (\rho_i, \rho_j)$, the sensitivity

$$m_j^\rho = \frac{1 + 2\rho_j + \rho_i(1 - \beta) + 2\rho_i\rho_j(1 - \beta^2)}{1 + 2(\rho_i + \rho_j) + 4\rho_i\rho_j(1 - \beta^2)}$$

of player j 's stage-2 equilibrium action to her primary signal is decreasing in ρ_i , implying that the sensitivity to her secondary signal, $1 - m_j^\rho$, is increasing in ρ_i . When actions are complements ($\beta > 0$) this in turn increases the incentives for player i to send player j a more precise secondary signal (so as to better coordinate with her), whereas the opposite is true when actions are substitutes ($\beta < 0$).

Next, consider ID. Fix the precision ρ_i of the secondary signal that player j expects to receive from player i . When player j sends a more precise signal to player i (that is, when $\hat{\rho}_j \geq \rho_j$), in the stage-2 game she then relies more on her primary signal if actions are complements ($\beta > 0$), and more on her secondary signal if actions are substitutes ($\beta < 0$). As a result, no matter the sign of

β , the incentives for player i to send player j a more precise secondary signal are smaller, implying that this game satisfies an inverted form of increasing differences. Combined, the above properties suggest that, in the case of strategic complements, such forms of manipulative cognition are likely to favor equilibria with asymmetric cognitive postures.

6 Endogenous depth of reasoning

The analysis in the preceding sections assumes that the players are fully rational, and that cognition takes the form of learning about payoffs and/or about other players' beliefs. In this section, we consider an alternative situation in which payoffs are common knowledge, but where players are boundedly rational and where cognition determines the players' ability to compute iterated best-responses. The analysis builds on the celebrated level- k model, in which k is a player's depth of reasoning, that is, the maximal number of steps of iterated best responses performed by the player. Contrary to the earlier literature (see, e.g., Crawford, Costa-Gomes, and Iriberry (2013) for a detailed overview), a player's depth of reasoning is endogenous. Alaoui and Penta (2016, 2017, 2018) are the first to endogenize the depth of reasoning in the level- k model. The main difference relative to their analysis is that we allow the value of expanding cognition to depend on a player's beliefs about (a) her opponents' cognition and (b) her opponents' expectations of her own cognitive level.

6.1 The environment

Consider the following two-player game in which payoffs are common knowledge. For each $i = 1, 2$, and each $k \in \mathbb{N}$, there is a mixed action $\alpha_i^k \in \Delta(A_i)$ such that α_i^k is a best response for player i to player j playing according to α_j^{k-1} , with α_i^0 specified exogenously, but reflecting a natural "anchor" that depends on the stage-2 game under consideration (up to this point, the formalism is the same as in the original model, where k is exogenous). Each player's cognitive level $\rho_i \in \mathbb{N}$ determines the player's endogenous depth of reasoning, that is, the number of steps of iterated best responses performed by the player. A player with depth of reasoning ρ_i who expects his opponent to have performed $\rho_j \geq \rho_i - 1$ steps of iterated best responses plays $\alpha_i^{\rho_i}$ in the stage-2 game. A player with depth of reasoning ρ_i who, instead, expects his opponent to have performed $\rho_j < \rho_i - 1$ steps of iterated best responses plays $\alpha_i^{\rho_j+1}$ in the stage-2 game. Formally, for any cognitive profile $\rho = (\rho_i, \rho_j)$, and any ρ'_i , player i 's period-2 mixed action when the two players are expected to choose cognition ρ and, instead, player i chooses cognition ρ'_i is given by

$$\sigma_i^{\rho'_i; \rho} = \begin{cases} \alpha_i^{\rho'_i} & \text{if } \rho'_i \leq \min\{\rho_i + 1, \rho_j\} + 1 \\ \alpha_i^{\min\{\rho_i+1, \rho_j\}+1} & \text{if } \rho'_i > \min\{\rho_i + 1, \rho_j\} + 1. \end{cases}$$

The idea is that player i plays the action corresponding to his cognitive capacity, unless, given the player's beliefs over the two players' cognitive capacities, player i believes his cognitive capacity exceeds the level that is necessary to perfectly predict the opponent's mixed action. This modeling of the stage-2 behavior is the same as in Alaoui and Penta (2016, 2017, 2018). As anticipated above, the key point of departure is in how players choose ρ_i . In Alaoui and Penta (2016, 2017, 2018), the choice of ρ_i is determined by a cost-benefit analysis in which both the costs and the benefits do not

depend on a player’s beliefs about her opponents’ cognitive choices and about their expectations of player i ’s own cognitive capacity. Here, instead, we allow for such dependence and investigate its implications for the selection of the cognitive levels.

Consistently with the notation used so far, continue to denote by

$$V_i(\rho'_i; \rho) = U_i(\sigma_i^{\rho'_i; \rho}, \sigma_j^\rho)$$

the ex-ante expected gross value of choosing cognition ρ'_i when the two players are expected to choose cognition $\rho = (\rho_i, \rho_j)$ and, instead, player i chooses cognition ρ'_i . Contrary to the rest of the analysis so far, though, V_i need not coincide with player i ’s value function. This naturally reflects the limited cognitive ability of the players (recall that this model is meant to be a description of the strategic reasoning of boundedly rational agents). Also note that we dropped ω from the player’s payoff function because, as explained above, in this game, payoffs are common knowledge (i.e., $|\Omega| = 1$).

In the spirit of Alaoui and Penta (2016, 2017, 2018), also assume that, when it comes to choosing their depth of reasoning, the players correctly perceive the value of V_i , even if they are not able to determine their correct best responses. Importantly, note that, contrary to the various cases considered in the previous sections, in this cognitive game, a player understands that, by increasing her cognition, she may end up with a lower payoff. This may happen despite the players’ cognitive choices being covert. The reason is that a player who increases her depth of reasoning but not to the point of being able to correctly identify the opponent’s true mixed action may find herself trapped into a cognitive loop that induces her to select a stage-2 mixed action that is farther away from her true best response than the one identified by computing a smaller number of iterated best responses.³²

6.2 Discussion

As mentioned above, the players correctly understand how their gross payoffs V_i depend on their own cognition, their opponent’s cognition, and their opponent’s expectation about their own cognition. This may feel at odds with the maintained assumption that the players need not be able to iterate their best responses enough to identify the rationalizable actions. It may also look unconvincing that players be able to predict how their stage-2 actions depend on their actual cognition, on their opponent’s cognition, and on their opponent’s expectations about their own cognition, without however being able to compute their precise best responses. These concerns are normal in models of bounded rationality. We suggest the reader does not interpret the model too literally. The model is meant to capture *forces* that may arise in the choice of the depth of reasoning. It seems quite plausible that such a choice depends on a player’s expectations of her opponent’s sophistication, as well as on her beliefs about her opponent’s expectation of her own sophistication. That a player’s perceived (gross) payoff $V_i(\rho'_i; \rho)$ from choosing cognition ρ'_i when the two players are expected to choose cognition ρ_i and ρ_j correctly reflects the dependence of the stage-2 actual actions on the players’ cognitive levels (and hence coincides with the player’s true payoff) is not essential. What matters is that a player correctly anticipates the forces that shape her stage-2 action, how they depend on the two

³²Clearly, on path, a player whose depth of knowledge exceeds her opponents’ never experiences a lower payoff. That higher cognition may backfire in the level-k model applies only to off-path cognitive choices.

players’ actual and perceived depth of reasoning, and how such forces in turn affect her payoff. In particular, what matters is that a player understands that (a) more cognition need not translate into higher payoffs when insufficient to identify the opponent’s action (it can even backfire by bringing a player’s action more farther apart from her true best response), (b) going significantly deeper into the understanding of the game than the opponent need not bring any advantage relative to going slightly deeper, and (c) once at her cognitive capacity, a player is unable to respond to variations in her opponent’s behavior due to a deeper understanding of the game. These features seem plausible and extend beyond the specific formalization above.³³

6.3 Arad and Rubinstein “11–20” game

For concreteness, we illustrate the role of expectation conformity in a specific game that has received considerable attention in the level- k literature. The stage-2 game described below was first introduced in Arad and Rubinstein (2012), and then simplified by Alaoui and Penta (2016). The players simultaneously announce an integer between 11 and 20. The players receive a number of tokens equal to the integer they announce. However, if a player announces an integer equal to the one announced by her opponent minus one, she receives extra x tokens, where $x \geq 20$. If the two players announce the same integer, they receive 10 tokens in addition to the integer they announce. Each token corresponds to one unit of payoff. Letting $A_i = \{11, 12, \dots, 20\}$, $i = 1, 2$, we thus have that the ex-post payoffs are equal to

$$u_i(a_i, a_j) = \begin{cases} a_i + x & \text{if } a_i = a_j - 1 \\ a_i + 10 & \text{if } a_i = a_j \\ a_i & \text{otherwise.} \end{cases}$$

This game, which is intended for experimental work, captures, in a stark and simplified manner, some of the forces that arise in certain strategic situations where players benefit from matching, or undercutting by little, the rivals’ actions. For example, the two players could be firms selling imperfectly substitutable goods to different segments of the market. If firm i ’s price exceeds the rival’s, firm i sells only to those consumers who do not value the rival’s product (the “loyalists”). Firm i is a monopolist on this segment of the market and its monopolistic price on this segment is $a_i = 20$. Reducing the price below $a_i = 20$ without attracting consumers who see the two goods as substitutes comes with a loss of profits. When, instead, firm i matches its rival’s price, in addition to selling to its loyalists, it also sells to 1/2 of those consumers who see the two products as substitutes. If it undercuts its rival, it sells to all consumers who see the two products as substitutes. However, if it undercuts the rival by a lot, the extra profits from conquering the contestable buyers are less than the losses from the loyalists. The Arad and Rubinstein game is meant to be a (highly simplified) version of the strategic situation that firms face in such circumstances.

Following Arad and Rubinstein (2012) and Alaoui and Penta (2016), let the “anchor” α_i^0 be the degenerate mixed action that selects the largest integer 20 with probability one.³⁴ This action is often interpreted as the most natural one in the absence of strategic reasoning. The version

³³As the reader may have recognized, some of these properties are in common with the sparsity model of Section 4.1.

³⁴Consistently with the analysis in the previous sections, $\delta_i(a_i)$ is the Dirac measure assigning probability one to the action $a_i \in A_i$.

originally introduced in Arad and Rubinstein (2012) does not feature the bonus of 10 tokens in case the players announce the same integer. The simplification proposed by Alaoui and Penta (2016) has two advantages: (a) it implies that, if the game was played by fully rational players, the unique rationalizable action would have both players select $a_i = 11$ with certainty; (b) it also implies that, for all i , and all $k \geq 9$, α_i^k is the degenerate mixed action that selects the integer 11 with certainty; that is, iterated best responses converge to the unique rationalizable action after 9 rounds. Because of this property, we simplify the analysis by assuming that $\rho_i \in \{0, 1, \dots, 9\}$, $i = 1, 2$. We then have the following result (the proof follows directly from the arguments after the proposition):

Proposition 12. (a) Consider any pair of cognitive profiles $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ and $\rho = (\rho_1, \rho_2)$ such that $\hat{\rho}_1 > \rho_1$, $\rho_2 = \rho_1 + 1$, and $\hat{\rho}_2 = \hat{\rho}_1 + 1$. Then $\Gamma_1^{UEC}(\rho, \hat{\rho}) = \Gamma_2^{UEC}(\rho, \hat{\rho}) = 0$, whereas $\Gamma_1^{ID}(\rho, \hat{\rho}) < 0 < \Gamma_2^{ID}(\rho, \hat{\rho})$. (b) Next, let $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ and $\rho = (\rho_1, \rho_2)$ be such that $\hat{\rho}_2 = \hat{\rho}_1 > \rho_2 = \rho_1$. Then, for $i = 1, 2$, $\Gamma_i^{UEC}(\rho, \hat{\rho}) = 0$ whereas $\Gamma_i^{ID}(\rho, \hat{\rho}) < 0$.

Hence, this game features a negative form of expectation conformity, at least with respect to the cognitive profiles under consideration. Consider first case (a). The idea behind this specific pair of cognitive profiles is the following. Suppose the two players are known to have different cognitive costs, with player 2 being the “leader” (that is, the player with the lowest cognitive cost). Further assume that both players’ cognitive costs are strictly increasing in their cognition. Then in any equilibrium in which the follower’s cognition is equal to ρ_1 , the leader’s cognition is equal to $\rho_2 = \rho_1 + 1$. Similarly, in any equilibrium in which the follower’s cognition is equal to $\hat{\rho}_1 > \rho_1$, the leader’s cognition is equal to $\hat{\rho}_2 = \hat{\rho}_1 + 1$.³⁵ We are interested in whether *multiple asymmetric equilibria* are possible in such a situation, driven by expectation conformity. The answer is no. To see why this is the case, consider first the situation faced by the follower (player 1). Fixing player 2’s cognitive level to $\rho_2 = \rho_1 + 1$, the value to player 1 of expanding her cognition from ρ_1 to $\hat{\rho}_1 \geq \rho_1 + 1 = \rho_2$ is (weakly) smaller when player 2 expects player 1 to choose cognition $\hat{\rho}_1$ than when she expects her to choose cognition $\rho_1 < \hat{\rho}_1$. To see this, note that the gross value to player 1 of expanding cognition from ρ_1 to $\hat{\rho}_1$ when player 2 expects player 1 to choose cognition $\hat{\rho}_1 \geq \rho_1 + 1 = \rho_2$ is equal to

$$\begin{aligned} V_1(\hat{\rho}_1; (\hat{\rho}_1, \rho_2)) - V_1(\rho_1; (\hat{\rho}_1, \rho_2)) &= [20 - \rho_2 + 10\mathbb{I}(\hat{\rho}_1 = \rho_1 + 1) + (x - 1)\mathbb{I}(\hat{\rho}_1 > \rho_1 + 1)] - (20 - \rho_1) \\ &= 10\mathbb{I}(\hat{\rho}_1 = \rho_1 + 1) + (x - 1)\mathbb{I}(\hat{\rho}_1 > \rho_1 + 1) - (\rho_2 - \rho_1). \end{aligned}$$

This is because player 2 announces $a_2 = 20 - \rho_2$ when she chooses cognition $\rho_2 = \rho_1 + 1$ and expects player 1 to choose cognition $\hat{\rho}_1 \geq \rho_1 + 1 = \rho_2$, whereas player 1, when she chooses cognition $\hat{\rho}_1 \geq \rho_1 + 1 = \rho_2$ and expects player 2 to choose cognition $\rho_2 = \rho_1 + 1$, she announces $a_1 = 20 - \rho_2$ if $\hat{\rho}_1 = \rho_1 + 1 = \rho_2$ and $a_1 = 20 - \rho_2 - 1$ if $\hat{\rho}_1 > \rho_1 + 1 = \rho_2$. When, instead, player 1 chooses cognition ρ_1 and expects player 2 to choose cognition $\rho_2 = \rho_1 + 1$, she then announces $a_1 = 20 - \rho_1$.

Likewise, when player 2 expects player 1 to choose cognition $\rho_1 = \rho_2 - 1$, she announces $a_2 = 20 - \rho_1 - 1 = 20 - \rho_2$. It follows that the gross value to player 1 of increasing her cognition from ρ_1 to $\hat{\rho}_1$ when player 2 expects her to choose cognition ρ_1 is the same as when player 2 expects her to choose cognition $\hat{\rho}_1$, implying that

$$\Gamma_1^{UEC}(\rho, \hat{\rho}) \equiv [V_1(\hat{\rho}_1; (\hat{\rho}_1, \rho_2)) - V_1(\rho_1; (\hat{\rho}_1, \rho_2))] - [V_1(\hat{\rho}_1; (\rho_1, \rho_2)) - V_1(\rho_1; (\rho_1, \rho_2))] = 0.$$

³⁵It is easy to verify that there exist no equilibrium in which the follower’s cognition is strictly higher than the leader’s.

Next, consider player 2 (the leader) and fix player 1's cognition to be equal to ρ_1 . The gross value to player 2 of expanding her cognition from $\rho_2 = \rho_1 + 1$ to $\hat{\rho}_2 = \hat{\rho}_1 + 1$ is the same no matter whether player 1 expects her to choose cognition ρ_2 or $\hat{\rho}_2$. This is because, in either case, player 1's ability to predict player 2's action is bounded by player 1's own cognitive capacity. In fact, player 1 announces $a_1 = 20 - \rho_1$, that is the action identified by ρ_1 steps of iterated best responses, irrespective of how far ahead she thinks player 2 is in the understanding of the game. This last property, which is the same as is Alaoui and Penta (2016), is similar in spirit to the one discussed in the context of sparsity in games. Hence, for player 2 as well, $\Gamma_2^{UEC}(\rho, \hat{\rho}) = 0$.

Next, consider ID, focusing again on the cognitive profiles of part (a) in the proposition. For player 1, the gross value of expanding her cognition from ρ_1 to $\hat{\rho}_1$ when player 2 expects her to choose $\hat{\rho}_1 \geq \rho_1 + 1 = \rho_2$ and chooses cognition $\hat{\rho}_2 = \hat{\rho}_1 + 1$ is equal to

$$V_1(\hat{\rho}_1; (\hat{\rho}_1, \hat{\rho}_2)) - V_1(\rho_1; (\hat{\rho}_1, \hat{\rho}_2)) = (20 - \hat{\rho}_1) - (20 - \rho_1) = -(\hat{\rho}_1 - \rho_1).$$

This is because, in this case, player 2 announces $a_2 = 20 - \hat{\rho}_1 - 1$, whereas player 1 announces $a_1 = 20 - \hat{\rho}_1$ when choosing cognition $\hat{\rho}_1$ and $a_1 = 20 - \rho_1$ when choosing cognition ρ_1 . The increase in cognition thus induces player 1 to lower her announcement, without, however, matching player 2's announcement, or undercutting it by one.

When, instead, player 2 expects player 1 to choose cognition $\hat{\rho}_1 \geq \rho_1 + 1 = \rho_2$ and chooses cognition ρ_2 , she then announces $a_2 = 20 - \rho_2$, in which case the value to player 1 of expanding her cognition from ρ_1 to $\hat{\rho}_1$ is equal to

$$V_1(\hat{\rho}_1; (\hat{\rho}_1, \rho_2)) - V_1(\rho_1; (\hat{\rho}_1, \rho_2)) = [20 - \rho_2 + 10\mathbb{I}(\hat{\rho}_1 = \rho_1 + 1) + (x - 1)\mathbb{I}(\hat{\rho}_1 > \rho_1 + 1)] - (20 - \rho_1).$$

Hence,

$$\begin{aligned} \Gamma_1^{ID}(\rho, \hat{\rho}) &\equiv [V_1(\hat{\rho}_1; (\hat{\rho}_1, \hat{\rho}_2)) - V_1(\rho_1; (\hat{\rho}_1, \hat{\rho}_2))] - [V_1(\hat{\rho}_1; (\hat{\rho}_1, \rho_2)) - V_1(\rho_1; (\hat{\rho}_1, \rho_2))] \\ &= -(\hat{\rho}_1 - \rho_1) - [10\mathbb{I}(\hat{\rho}_1 = \rho_1 + 1) + (x - 1)\mathbb{I}(\hat{\rho}_1 > \rho_1 + 1) - (\rho_2 - \rho_1)] \\ &= -(\hat{\rho}_1 - \rho_2) - [10\mathbb{I}(\hat{\rho}_1 = \rho_1 + 1) + (x - 1)\mathbb{I}(\hat{\rho}_1 > \rho_1 + 1)] < 0. \end{aligned}$$

The reason why this game features a negative form of increasing differences for the player with the highest cognitive cost (equivalently, with the lowest expected cognitive level) is that player 2, when expecting player 1's cognition to be equal to $\hat{\rho}_1 \geq \rho_1 + 1 = \rho_2$, announces $a_2 = 20 - \hat{\rho}_1 - 1$ when choosing the high cognitive level $\hat{\rho}_2 = \hat{\rho}_1 + 1$, whereas she announces $a_2 = 20 - \rho_1 + 1$ when choosing the low cognitive level $\rho_2 = \rho_1 + 1$. Thus player 1 (the one who is expected to be behind in the exploration of the game) suffers from an increase in cognition by her opponent.

Next consider player 2 (the one who is expected to be ahead in the exploration of the game). When player 1 expects her to choose cognition $\hat{\rho}_2$, the gross value of increasing her cognition from $\rho_2 = \rho_1 + 1$ to $\hat{\rho}_2 = \hat{\rho}_1 + 1$ is equal to

$$V_2(\hat{\rho}_2; (\hat{\rho}_1, \hat{\rho}_2)) - V_2(\rho_2; (\hat{\rho}_1, \hat{\rho}_2)) = 20 - \hat{\rho}_2 + x - [(20 - \rho_2) + 10\mathbb{I}(\hat{\rho}_1 = \rho_1 + 1)] > 0$$

when player 1's cognition is equal to the high level $\hat{\rho}_1$ and is equal to

$$V_2(\hat{\rho}_2; (\rho_1, \hat{\rho}_2)) - V_1(\rho_2; (\rho_1, \hat{\rho}_2)) = 0$$

when player 1's cognition is equal to the low level ρ_1 (In this latter case, player 2 expects her capacity to be large enough to perfectly predict player 1's announcement, no matter whether she chooses $\rho_2 = \rho_1 + 1$ or $\hat{\rho}_2 = \hat{\rho}_1 + 1 > \rho_2$). It follows that, for the leader, this game features positive increasing differences:

$$\Gamma_2^{ID}(\rho, \hat{\rho}) \equiv [V_2(\hat{\rho}_2; (\hat{\rho}_1, \hat{\rho}_2)) - V_2(\rho_2; (\hat{\rho}_1, \hat{\rho}_2))] - [V_2(\hat{\rho}_2; (\rho_1, \hat{\rho}_2)) - V_1(\rho_2; (\rho_1, \hat{\rho}_2))] > 0.$$

Combining the results for unilateral expectation conformity with those for increasing differences, we conclude that, in this game, $\Gamma_1^{EC}(\rho, \hat{\rho}) < 0 < \Gamma_2^{EC}(\rho, \hat{\rho})$. Arguments similar to those establishing Proposition 1 then imply that, when the cognitive costs are strictly increasing, there cannot exist multiple asymmetric equilibria.

Next, consider the cognitive profiles in part (b) of the proposition. What motivates considering such profiles is their relation to the possibility of *multiple symmetric equilibria*. The result in the proposition implies that such a multiplicity is not possible.

First, consider UEC and, without loss of generality, focus on player 2. When player 1 chooses cognition ρ_1 , in the stage-2 game, she then announces $a_1 = 20 - \rho_1$, no matter whether she expects player 2 to choose $\rho_2 = \rho_1$ or $\hat{\rho}_2 > \rho_2 = \rho_1$. This is because, in either case, player 1 is at her cognitive capacity. As a consequence, the value to player 2 of increasing her cognition from ρ_2 to $\hat{\rho}_2 > \rho_2$ is positive but invariant to player 1's expectations about player 1's cognition. From the definition of $\Gamma_2^{UEC}(\rho, \hat{\rho})$, we then have that

$$\begin{aligned} \Gamma_2^{UEC}(\rho, \hat{\rho}) &\equiv [V_2(\hat{\rho}_2; (\rho_1, \hat{\rho}_2)) - V_2(\rho_2; (\rho_1, \hat{\rho}_2))] - [V_2(\hat{\rho}_2; (\rho_1, \rho_2)) - V_2(\rho_2; (\rho_1, \rho_2))] \\ &= (20 - \rho_1 - 1 + x) - (20 - \rho_1) - [(20 - \rho_1 - 1 + x) - (20 - \rho_1)] = 0. \end{aligned}$$

Because the two players face the same situation under the cognitive profiles under consideration, the same conclusion applies to player 1, that is, $\Gamma_1^{UEC}(\rho, \hat{\rho}) = 0$.

Next, consider ID. When player 1 expects player 2 to choose a higher cognitive level $\hat{\rho}_2$, she then announces $a_1 = 20 - \rho_1$ when choosing the low cognitive level $\rho_1 = \rho_2$ and $a_1 = 20 - \hat{\rho}_1$ when choosing the high cognitive level $\hat{\rho}_1 = \hat{\rho}_2$ (in both cases, player 1 is constrained by her cognitive capacity). The value to player 2 of increasing her cognition from ρ_2 to $\hat{\rho}_2$ is then equal to

$$V_2(\hat{\rho}_2; (\hat{\rho}_1, \hat{\rho}_2)) - V_2(\rho_2; (\hat{\rho}_1, \hat{\rho}_2)) = 20 - \hat{\rho}_2 + 10 - (20 - \rho_2)$$

when player 1 chooses the high cognitive level $\hat{\rho}_1 = \hat{\rho}_2$, whereas it is equal to

$$V_2(\hat{\rho}_2; (\rho_1, \hat{\rho}_2)) - V_2(\rho_2; (\rho_1, \hat{\rho}_2)) = 20 - \rho_2 - 1 + x - (20 - \rho_2 + 10)$$

when player 1 chooses the low cognitive level $\rho_1 = \rho_2$. Because $x > 20$, we then have that

$$\Gamma_2^{ID}(\rho, \hat{\rho}) \equiv [V_2(\hat{\rho}_2; (\hat{\rho}_1, \hat{\rho}_2)) - V_2(\rho_2; (\hat{\rho}_1, \hat{\rho}_2))] - [V_2(\hat{\rho}_2; (\rho_1, \hat{\rho}_2)) - V_2(\rho_2; (\rho_1, \hat{\rho}_2))] < 0.$$

The same conclusion applies to player 1. This game thus features a form of negative ID with respect to the profiles under consideration: increasing cognition is more valuable when the opponent chooses a lower cognition. Again, arguments similar to those in Proposition 1 then imply that this game, despite being played by boundedly rational players, cannot feature multiple symmetric (pure-strategy) equilibria.

7 Concluding remarks

Before playing a game, economic agents think about the strategic situation they are facing: they acquire information about payoffs, engage in brainstorming with other agents (financial experts, consultants, but also individuals who faced similar strategic situations in the past), invest in making themselves understood (for example, by communicating their views), “spy” on other agents to learn other agents’ view of the game, but also iterate over the computation of the best responses to identify the actions that are most likely to maximize their payoffs.

Understanding how players make such cognitive investments may help predict the way specific games are played. It can also help interpreting the functioning of contracts, the (in)efficiency of certain markets, and the way social interactions can be influenced by policy interventions. These observations motivate the study of “cognitive games,” that is, strategic situations where players engage in various cognitive activities shaping their own understanding of the game, as well as others’.

In this paper, we focus on a specific aspect of such games: the role of expectations. We introduce a notion of expectation conformity which we then use to deliver predictions about the selection of the cognitive structures. We show how expectation conformity is driven by the interaction between two synergies, (a) the value to conform to other players’ expectations about one’s own cognition (unilateral expectation conformity), and (b) the value to conform to other players’ actual cognitive choices (increasing differences).

We investigate how expectation conformity (or the lack thereof) relates to the nature of the primitive game (in particular, whether actions are strategic complements or substitutes) and how it contributes to the determinacy of equilibria and the possibility of cognitive traps. We find that constant-sum games never give rise to self-fulfilling cognition, whereas linear-quadratic games often do (both when actions are complements and when they are substitutes). We consider both the case in which cognition is self-directed (i.e., it affects a player’s understanding of the game without affecting others’) as well as the case where cognition is manipulative (i.e., it affects other players’ understanding of the game). The analysis is applied to different cognitive modes, including sparsity, noisy information acquisition about exogenous payoff states, espionage, framing and memory management. Finally, we investigate how expectations shape the depth of reasoning in models of bounded rationality such as level-k thinking.

Throughout the analysis, we confine attention to cognitive investments that are essentially static in nature: players make such investments once and for all prior to playing the primitive game. We also confine attention to situations where such investments are covert, that is, not directly observable by other players. In future work, it would be interesting to extend the analysis to settings in which agents adjust their cognition as the game progresses. This possibility introduces new effects, such as the possibility that players sacrifice current flow payoffs to enhance their understating of the game

and exploit the acquired knowledge in subsequent periods (as in the multi-armed bandit literature). It also introduces the possibility for players to signal their cognition to other players, a dimension that appears relevant in certain problems of interest.

Throughout the paper, we also confine attention to positive analysis. We do not discuss any normative aspect related to the selection of the cognitive structures. In future work, it would be interesting to investigate whether players over- or under-invest in cognition and how the latter properties depend on the type of strategic situation and on the form of cognition. It would also be interesting to relate inefficiencies in cognition to inefficiencies in the way the primitive game is played, and use the analysis to identify interventions that can mitigate such inefficiencies.

References

- Adriani, F., and S. Sonderegger (2020), "Listening to the opinions of others: attention costs and asymmetric attention patterns," mimeo Leicester and Nottingham Universities.
- Alaoui, L. and A. Penta (2016), "Endogenous Depth of Reasoning," *Review of Economic Studies*, 83, 1297–1333.
- Alaoui, L. and A. Penta (2016), "Reasoning about Others' Reasoning," mimeo, University Pompeu Fabra.
- Alaoui, L. and A. Penta (2018), "Cost-Benefit Analysis in Reasoning," mimeo, University Pompeu Fabra.
- Amir, R., and N. Lazzati (2016) "Endogenous Information Acquisition in Bayesian Games with Strategic Complementarities," *Journal of Economic Theory*, 163: 684–698.
- Angeletos, G.M., and A. Pavan (2007), "Efficient Use of Information and Social Value of Information," *Econometrica* 75: 1103–1142.
- Angeletos G.M. and K. Sastry (2020), "Inattentive Economies," mimeo MIT and Yale university.
- Arad, A. and A. Rubinstein (2012), "The 11-20 Money Request Game: A Level-k Reasoning Study," *American Economic Review*, 102 (7): 3561-73.
- Banerjee, S., Davis, J. and N. Gondhi, (2020), "Motivated Beliefs in Coordination Games," mimeo UNC.
- Battigalli, P. and M. Dufwenberg, (2009), "Dynamic psychological games," *Journal of Economic Theory* 144(1), 1-35.
- Bénabou, R. (2013) "Groupthink: Collective Delusions in Organizations and Markets," *Review of Economics Studies*, 80: 429–462.
- Bénabou, R., and J. Tirole (2002) "Self-Confidence and Personal Motivation," *Quarterly Journal of Economics*, 117(3): 871–915.
- Bergemann, D. and J. Välimäki (2002) "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 70(3), 1007–1033.
- Bergemann, D., and S. Morris (2013) "Robust Predictions in Games with Incomplete Information" *Econometrica*, 81(4), 1251-1308.
- Bergemann, D., and S. Morris (2019) "Information Design: A Unified Perspective," *Journal of Economic Literature*, 57, 44–95.
- Bolton, P., and A. Faure-Grimaud (2010) "Satisficing Contracts," *Review of Economic Studies*, 77(3): 937–971.
- Calvó-Armengol, A. J. De Martí, and A. Prat, (2015), "Communication and influence," *Theoretical Economics* 10.2 649-690.
- Caplin, A., M. Dean, and J. Leahy, (2018), "Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy," mimeo Columbia, NYU, and University of Michigan.
- Che, Y.K., and N. Kartik (2009) "Opinions as Incentives," *Journal of Political Economy*, 117(5): 815–860.
- Colombo, I., Femminis, G., and A. Pavan (2014) "Information Acquisition and Welfare," *Review of Economic Studies*, 81: 1438–1483.

- Crémer, J., and F. Khalil (1992) "Gathering Information Before Signing a Contract," *American Economic Review*, 82(3): 566–578.
- Crémer, J., and F. Khalil (1994) "Gathering Information before the Contract is Offered: the Case with Two States of Nature," *European Economic Review*, 38(3-4): 675–682.
- Crémer, J., Khalil, F., and J.C. Rochet (1998a) "Contracts and Productive Information Gathering." *Games and Economic Behavior*, 25(2): 174–193.
- Crémer, J., Khalil, F., and J.C. Rochet (1998b) "Strategic Information Gathering before a Contract is Offered." *Journal of Economic Theory*, 81(1): 163–200.
- Crawford, V., M. Costa-Gomes, N. Iriberri (2013) "Structural Models of Non-equilibrium Strategic Thinking: Theory, Evidence, and Applications," *Journal of Economic Literature*, 51 (1): 5-62.
- Dang, T.V. (2008) "Bargaining with Endogenous Information," *Journal of Economic Theory*, 140: 339–354.
- Dang, T.V., Gorton, G., Holmström, B. and G. Ordonez (2017) "Banks as Secret Keepers," *American Economic Review*, 107: 1005–1029.
- Dessi, R. (2008) "Collective Memory, Cultural Transmission and Investments," *American Economic Review*, 98(1): 534–560.
- Denti, T. (2020), "Unrestricted Information Acquisition," mimeo Cornell University.
- Dewatripont, M., Jewitt, I., and J. Tirole (1999) "The Economics of Career Concerns. I: Comparison of Information Structures," *Review of Economic Studies*, 66 (1): 183–198.
- Dewatripont, M., and J. Tirole (2005) "Modes of Communication," *Journal of Political Economy*, 113: 1217–1238.
- Farhi, E., and J. Tirole (2015) "Liquid Bundles," *Journal of Economic Theory*, 158(B): 634–655.
- Fudenberg, D., and J. Tirole (1986), "A signal-jamming theory of predation," *RAND Journal of Economics*: 366-376.
- Gabaix, X. (2014) "A Sparsity-Based Model of Bounded Rationality," *Quarterly Journal of Economics*, 129(4): 1661–1710.
- Geanakoplos, J., D. Pearce, and E. Stacchetti, (1989) "Psychological Games and Sequential Rationality," *Games and Economic Behavior*, 60–79.
- Gottlieb, D. (2014a) "Imperfect Memory and Choice under Risk," *Games and Economic Behavior*, 85: 127–185.
- Gottlieb, D. (2014b) "Will you Never Learn? Self-Deception and Biases in Information Processing," mimeo, Wharton School, University of Pennsylvania.
- Hebert, B. and J. La'O (2021) "Information Acquisition, Efficiency, and Non-Fundamental Volatility," NBER Working Paper No. 26771.
- Hebert, B. and M. Woodford (2020), "Neighborhood-Based Information Costs," NBER Working Paper 26743.
- Hellwig, C., and L. Veldkamp (2009) "Knowing What Others Know: Coordination Motives in Information Acquisition," *Review of Economic Studies*, 76(1): 223–251.
- Holmström B. (1999) "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1): 169–182.

- Horner, J. and N. Lambert, (2019), “Motivational Ratings,” mimeo Yale and Stanford Universities.
- Kamenica, E., (2019), “Bayesian Persuasion and Information Design,” *Annual Review of Economics*, forthcoming.
- Kozlovskaya, M. (2018) “Industrial Espionage in Duopoly Games,” mimeo, University of Aston.
- Lehrer, E., and D. Rosenberg (2006) “What Restrictions do Bayesian Games Impose on the Value of Information?,” *Journal of Mathematical Economics*, 42: 343–357.
- Liang, A., and X. Mu (2020) “Complementary Information and Learning Traps,” *Quarterly Journal of Economics*, 135(1): 389–448.
- Maćkowiak, B., F. Matejka, and M. Wiederholt, (2020) “Rational Inattention: A Review,” mimeo CERGE-EI.
- Maćkowiak, B., and M. Wiederholt (2009) “Optimal Sticky Prices under Rational Inattention,” *American Economic Review*, 99(3): 769–803.
- Matejka, F., and A. McKay (2012) “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model,” *American Economic Review*, 105(1): 272–298.
- Matejka, F., J. Steiner, and C. Stewart (2017) “Rational Inattention Dynamics: Inertia and Delay in Decision-making,” *Econometrica* 85(2), 521–553.
- Morris, S., and H. Shin (2002) “Social Value of Public Information,” *American Economic Review*, 92: 1521–1534.
- Morris, S., and M. Yang (2019) “Coordination and Continuous Stochastic Choice,” mimeo MIT and Duke University.
- Mullainathan, S. (2002) “Thinking through Categories”, mimeo, MIT.
- Mullainathan, S., Schwartzstein, J. and A. Shleifer (2008) “Coarse Thinking and Persuasion,” *Quarterly Journal of Economics*, 123: 577–619.
- Myatt, D. P., and C. Wallace (2012) “Endogenous Information Acquisition in Coordination Games,” *Review of Economic Studies*, 79: 340–374.
- Pavan, A. (2016) “Attention, Coordination, and Bounded Recall,” mimeo, Northwestern University.
- Pavan, A., and J. Tirole (2021a) “Knowing your Lemon before you Dump It,” mimeo, Northwestern University and Toulouse School of Economics.
- Pavan, A, and J. Tirole, (2021b) “Duty of Disclosure and Contract Design,” mimeo, Toulouse School of Economics.
- Persico, N. (2000) “Information Acquisition in Auctions,” *Econometrica*, 68(1): 135–148.
- Salant, Y., and R. Siegel (2018), “Contracts with Framing,” *American Economic Journal: Microeconomics*, 10 (3): 315–46.
- Sims, C. (2003) “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50(3): 665–690.
- Sethi, R. and M. Yildiz (2016), “Communication with Unknown Perspectives,” *Econometrica*, 84(6) 2029–2069.
- Sethi, R. and M. Yildiz (2018), “Culture and Communication,” mimeo MIT.

- Stigler, G. (1961) “The Economics of Information,” *Journal of Political Economy*, 69: 213–225.
- Szkup, M. and I. Trevino (2015) “Information Acquisition in Global Games of Regime Change,” *Journal of Economic Theory* 160, 387-428.
- Tirole, J. (2009) “Cognition and Incomplete Contracts,” *American Economic Review*, 99(1): 265–294.
- Yang, M. (2015) “Coordination with Flexible Information Acquisition,” *Journal of Economic Theory*, 158: 721–738.
- Yang, M. (2020), “Optimality of Debt under Flexible Information Acquisition,” *Review of Economic Studies*, 87(1), 487-536.

Appendix

Proof of Proposition 4. We start by establishing the general properties of the individual best responses and of the ex-ante gross expected payoff functions V_i discussed in the paragraphs preceding the proposition. Equipped with the aforementioned results, we then establish the properties in the proposition.

First note that the players' payoffs can be rewritten as

$$u_i(a_i, a_j, \omega) = 2a_i [(1 - \beta)g(\omega) + \beta a_j] - a_i^2 + \{\psi(a_{-i}, \omega) - (1 - \beta)g(\omega)^2 - \beta \bar{a}_{-i}^2\}.$$

Hereafter, we disregard the terms in the curly bracket because these are external effects that have no impact on individual best responses and hence cancel out in the payoff differentials used to compute ID and UEC. Using the assumption that $g(\omega) = (1 + \bar{\omega})/(1 - \beta)$, we then proceed as if the agents' payoffs were given by

$$\hat{u}_i(a_i, a_{-i}, \omega) = 2a_i \left(1 + \sum_{k=1}^K \omega^k + \beta \bar{a}_{-i} \right) - a_i^2. \quad (19)$$

As it has been noticed in the literature, one can then interpret a_i as the agent's investment into a risky activity, 1 as the component of the gross return to such an investment that is commonly known, $\bar{\omega}$ as the sum of the components (to the gross return) that are unknown, $\beta \bar{a}_{-i}$ as an investment spillover (positive for $\beta > 0$ and negative for $\beta < 0$), and a_i^2 as a quadratic adjustment cost.

To further simplify the derivations, we also assume that there are only two players, that is, $n = 2$.

Best responses and equilibrium actions. Fix the cognitive profile $\rho = (\rho_1, \rho_2)$, with $\rho_1 \leq \rho_2$, and note that, given any ω , the two players' best responses must satisfy the optimality conditions

$$a_i(\omega_i) = 1 + \bar{\omega}_i + \beta \mathbb{E}[\sigma_{-i}(\omega_{-i}) | \omega_i, \rho]$$

for $i, j = 1, 2$, $j \neq i$, where $\bar{\omega}_i \equiv \sum_{k=1}^{\rho_i} \omega^k$. Because both players' stage-2 continuation payoffs are strictly quasi-concave in their own actions, the stage-2 equilibrium strategies σ_i must be Dirac measures assigning probability one to actions satisfying the above optimality conditions.

That player 2 has a larger knowledge than player 1 in turn implies that, for any ω_2 , $\mathbb{E}[\sigma_1(\omega_1) | \omega_2, \rho] = a_1(\omega_1)$. Player's 2 equilibrium action must thus satisfy the optimality condition

$$a_2(\omega_2) = 1 + \bar{\omega}_2 + \beta a_1(\omega_1)$$

for all ω_2 . In turn, this implies that, for any ω_1 , player 1's stage-2 equilibrium action must satisfy

$$a_1(\omega_1) = 1 + \bar{\omega}_1 + \beta [1 + \bar{\omega}_1 + \beta a_1(\omega_1)].$$

Combining the above two best responses, we then have that the stage-2 equilibrium actions are given by

$$a_1^\rho(\omega_1) = \frac{1 + \bar{\omega}_1}{1 - \beta}$$

and

$$a_2^\rho(\omega_2) = a_1^\rho(\omega_1) + \bar{\omega}_2 - \bar{\omega}_1.$$

In each state ω , the two players' equilibrium interim expected gross payoffs are then equal to

$$\mathbb{E}[\hat{u}_1(a_1^\rho(\omega_1), a_2^\rho(\omega_2), \omega)|\omega_1, \rho] = 2a_1^\rho(\omega_1)(1 + \bar{\omega}_1 + \beta a_1^\rho(\omega_1)) - a_1^\rho(\omega_1)^2 = \left(\frac{1 + \bar{\omega}_1}{1 - \beta}\right)^2$$

and

$$\mathbb{E}[\hat{u}_2(a_1^\rho(\omega_1), a_2^\rho(\omega_2), \omega)|\omega_2, \rho] = 2a_2^\rho(\omega_2)(1 + \bar{\omega}_2 + \beta a_1^\rho(\omega_1)) - a_2^\rho(\omega_2)^2 = \left(\frac{1 + \bar{\omega}_2 - \beta(\bar{\omega}_2 - \bar{\omega}_1)}{1 - \beta}\right)^2.$$

Integrating over states, we then have that the ex-ante expected gross payoffs are equal to

$$V_1(\rho_1; \rho) = \frac{1 + \sum_{k=1}^{\rho_1} \sigma_k^2}{(1 - \beta)^2}$$

and

$$V_2(\rho_2; \rho) = V_1(\rho_1; \rho) + \sum_{k=\rho_1+1}^{\rho_2} \sigma_k^2.$$

Given the derivations above, it is easy to verify that player 1's actions and payoff coincide with those that that would prevail in a fictitious environment in which (a) all dimensions ω_k , $k > \rho_1$, are equal to zero and such event is commonly known, and (b) all dimensions ω_k , $k \leq \rho_1$, are commonly known by the two players. It is also easy to verify that player 2's actions and payoff coincide with those that that would prevail in a fictitious environment in which (a) all dimensions $k > \rho_2$ are equal to zero and such event is commonly known, (b) all dimensions $k \leq \rho_1$ are commonly known, (c) it is commonly known that player 1 believes that all dimensions $k > \rho_1$ are identically equal to zero, and (d) player 2 knows all dimensions ω_k with $k \leq \rho_2$. In this sense, each player reasons, and acts, "as if" all dimensions $k > \rho_i$ simply "did not exist", as claimed in the main body.

Player 1's deviations. Now suppose that the two players are expected to choose cognition $\rho = (\rho_1, \rho_2)$, with $\rho_1 < \rho_2$, and player 1 deviates and selects cognition ρ'_1 with $\rho_1 < \rho'_1 \leq \rho_2$. Then, for any

$$\omega'_1 \equiv (\omega^1, \dots, \omega^{\rho_i}, \omega^{\rho_i+1}, \dots, \omega^{\rho'_i})$$

his optimal stage-2 strategy consists in choosing with probability one the action

$$\begin{aligned} a_1^{\rho'_1; \rho}(\omega'_1) &= 1 + \bar{\omega}'_1 + \beta \mathbb{E}[a_2(\omega_2)|\omega'_1, \rho'_1, \rho_2] \\ &= 1 + \bar{\omega}'_1 + \beta (a_1^\rho(\omega_1) + \bar{\omega}'_1 - \bar{\omega}_1) \\ &= a_1^\rho(\omega_1) + (1 + \beta)(\bar{\omega}'_1 - \bar{\omega}_1), \end{aligned}$$

where $\bar{\omega}_1 \equiv \sum_{k=1}^{\rho_1} \omega^k$, and $\bar{\omega}'_1 \equiv \sum_{k=1}^{\rho'_1} \omega^k$.

It follows that player 1's interim expected payoff following the deviation is equal to

$$\mathbb{E}[\hat{u}_1(a_1^{\rho'_1; \rho}(\omega'_1), a_2^\rho(\omega_2), \omega)|\omega'_1, (\rho'_1, \rho_2)] = 2a_1^{\rho'_1; \rho}(\omega'_1)(1 + \bar{\omega}'_1 + \beta [a_1^\rho(\omega_1) + \bar{\omega}'_1 - \bar{\omega}_1]) - a_1^{\rho'_1; \rho}(\omega'_1)^2$$

which simplifies to

$$\begin{aligned} \mathbb{E} \left[\hat{u}_1(a_1^{\rho'_1; \rho}(\omega'_1), a_2^\rho(\omega_2), \omega) | \omega'_1, (\rho'_1, \rho_2) \right] &= 2a_1^\rho(\omega_1) (1 + \bar{\omega}_1 + \beta a_1^\rho(\omega_1)) - a_1^\rho(\omega_1)^2 \\ &+ 2a_1^\rho(\omega_1)(1 + \beta)(\bar{\omega}'_1 - \bar{\omega}_1) + 2(1 + \beta)(\bar{\omega}'_1 - \bar{\omega}_1) (1 + \bar{\omega}_1 + \beta a_1^\rho(\omega_1)) \\ &+ 2(1 + \beta)^2(\bar{\omega}'_1 - \bar{\omega}_1)^2 - (1 + \beta)^2(\bar{\omega}'_1 - \bar{\omega}_1)^2 - 2a_1^\rho(\omega_1)(1 + \beta)(\bar{\omega}'_1 - \bar{\omega}_1). \end{aligned}$$

Integrating over states we then have that his ex-ante expected gross payoff following the deviation is equal to

$$V_1(\rho'_1; \rho) = V_1(\rho_1; \rho) + (1 + \beta)^2 \sum_{k=\rho_1+1}^{\rho'_1} \sigma_k^2.$$

Next, consider the case where the two players are expected to choose cognition $\rho = (\rho_1, \rho_2)$, with $\rho_1 \leq \rho_2$, and player 1 deviates and selects cognition ρ'_1 with $\rho'_1 > \rho_2$. Then, for any

$$\omega'_1 \equiv (\omega^1, \dots, \omega^{\rho_i}, \omega^{\rho_i+1}, \dots, \omega^{\rho'_i}),$$

his optimal stage-2 strategy consists in choosing with probability one the action

$$\begin{aligned} a_1^{\rho'_1; \rho}(\omega'_1) &= 1 + \bar{\omega}'_1 + \beta \mathbb{E} [a_2(\omega_2) | \omega'_1, (\rho'_1, \rho_2)] \\ &= 1 + \bar{\omega}_1 + (\bar{\omega}'_1 - \bar{\omega}_1) + \beta (a_1^\rho(\omega_1) + \bar{\omega}_2 - \bar{\omega}_1) \\ &= a_1^\rho(\omega_1) + (1 + \beta)(\bar{\omega}_2 - \bar{\omega}_1) + (\bar{\omega}'_1 - \bar{\omega}_2) \\ &= a_2^\rho(\omega_2) + \beta(\bar{\omega}_2 - \bar{\omega}_1) + (\bar{\omega}'_1 - \bar{\omega}_2). \end{aligned}$$

It follows that player 1's interim expected payoff following the deviation is equal to

$$\mathbb{E} \left[\hat{u}_1(a_1^{\rho'_1; \rho}(\omega'_1), a_2^\rho(\omega_2), \omega) | \omega'_1, (\rho'_1, \rho_2) \right] = 2a_1^{\rho'_1; \rho}(\omega'_1) (1 + \bar{\omega}'_1 + \beta a_2^\rho(\omega_2)) - a_1^{\rho'_1; \rho}(\omega'_1)^2$$

which simplifies to

$$\begin{aligned} \mathbb{E} \left[\hat{u}_1(a_1^{\rho'_1; \rho}(\omega'_1), a_2^\rho(\omega_2), \omega) | \omega'_1, (\rho'_1, \rho_2) \right] &= 2a_2^\rho(\omega_2) (1 + \bar{\omega}_2 + \beta a_1^\rho(\omega_1)) - a_2^\rho(\omega_2)^2 \\ &+ 2\beta(\bar{\omega}_2 - \bar{\omega}_1) (1 + \bar{\omega}_1 + \beta a_1^\rho(\omega_1)) + 2\beta(1 + \beta)(\bar{\omega}_2 - \bar{\omega}_1)^2 + 2\beta(\bar{\omega}_2 - \bar{\omega}_1)(\bar{\omega}'_1 - \bar{\omega}_2) \\ &+ 2(\bar{\omega}'_1 - \bar{\omega}_2) [1 + \bar{\omega}_2 + \beta (a_1^\rho(\omega_1) + \bar{\omega}_2 - \bar{\omega}_1)] + 2(\bar{\omega}'_1 - \bar{\omega}_2)^2 \\ &- \beta^2(\bar{\omega}_2 - \bar{\omega}_1)^2 - (\bar{\omega}'_1 - \bar{\omega}_2)^2 - 2\beta(\bar{\omega}_2 - \bar{\omega}_1)(\bar{\omega}'_1 - \bar{\omega}_2). \end{aligned}$$

Integrating over states we then have that his ex-ante expected gross payoff following the deviation is equal to

$$\begin{aligned} V_1(\rho'_1; \rho) &= V_2(\rho_2; \rho) + 2\beta \left(1 + \frac{\beta}{2} \right) \sum_{k=\rho_1+1}^{\rho_2} \sigma_k^2 + \sum_{k=\rho_2+1}^{\rho'_1} \sigma_k^2 \\ &= V_1(\rho_1; \rho) + (1 + \beta)^2 \sum_{k=\rho_1+1}^{\rho_2} \sigma_k^2 + \sum_{k=\rho_2+1}^{\rho'_1} \sigma_k^2. \end{aligned}$$

Finally, suppose that the two players are expected to choose cognition $\rho = (\rho_1, \rho_2)$, with $\rho_1 \leq \rho_2$, and that player 1 deviates to some lower cognition $\rho'_1 < \rho_1$. Because in this case his optimal action is equal to

$$\begin{aligned}
a_1^{\rho'_1; \rho}(\omega'_1) &= 1 + \bar{\omega}_1' + \beta \mathbb{E}[a_2(\omega_2) | \omega'_1, (\rho'_1, \rho_2)] \\
&= 1 + \bar{\omega}_1' + \beta \mathbb{E}[a_1^\rho(\omega_1) + \bar{\omega}_2 - \bar{\omega}_1 | \omega'_1, (\rho'_1, \rho_2)] \\
&= 1 + \bar{\omega}_1' + \beta \frac{1 + \bar{\omega}_1'}{1 - \beta} \\
&= a_1^{\rho'_1, \rho_2}(\omega_1).
\end{aligned}$$

it is evident that his ex-ante expected payoff is equal to

$$V_1(\rho'_1; \rho) = V_1(\rho'_1; (\rho'_1, \rho_2)) = \frac{1 + \sum_{k=1}^{\rho'_1} \sigma_k^2}{(1 - \beta)^2}.$$

Player 2's deviations. Next, consider deviations by the player who is supposed to be ahead in the exploration of the state. Suppose that the two players are expected to choose cognition $\rho = (\rho_1, \rho_2)$, with $\rho_1 \leq \rho_2$, and player 2 deviates to cognition ρ'_2 with $\rho'_2 \geq \rho_1$. His optimal stage-2 action is then equal to

$$a_2^{\rho'_2; \rho}(\omega'_2) = 1 + \bar{\omega}_2' + \beta \mathbb{E}[a_1^\rho(\omega_1) | \omega'_2, (\rho_1, \rho'_2)] = 1 + \bar{\omega}_2' + \beta a_1^\rho(\omega_1) = a_2^{\rho_1, \rho'_2}(\omega'_2).$$

That is, the optimal action for player 2 coincides with the action that she would choose in equilibrium if his deviation was observable. The ex-ante expected gross payoff that player 2 obtains following such deviation is thus equal to

$$V_2(\rho'_2; \rho) = V_2(\rho'_2; (\rho_1, \rho'_2)) = V_1(\rho_1; \rho) + \sum_{k=\rho_1+1}^{\rho'_2} \sigma_k^2.$$

Finally, consider the payoff that player 2 obtains by deviating to cognition $\rho'_2 < \rho_1$. His optimal stage-2 action is then equal to

$$a_2^{\rho'_2; \rho}(\omega'_2) = 1 + \bar{\omega}_2' + \beta \mathbb{E}[a_1^\rho(\omega_1) | \omega'_2, \rho_1, \rho'_2] = 1 + \bar{\omega}_2' + \beta \frac{1 + \bar{\omega}_2'}{1 - \beta} = \frac{1 + \bar{\omega}_2'}{1 - \beta} = a_2^{\rho_1, \rho'_2}(\omega'_2).$$

Again, player 2's optimal action coincides with the action that he would take if his deviation was observable. In turn, this also implies that his ex-ante expected gross payoff following such a deviation is equal to

$$V_2(\rho'_2; \rho) = V_2(\rho'_2; (\rho_1, \rho'_2)) = \frac{1 + \sum_{k=1}^{\rho'_2} \sigma_k^2}{(1 - \beta)^2}.$$

Expectation conformity.

We are now ready to establish the results in the proposition. To see whether UEC, consider first the case where the player who is expected to be ahead is the same under all relevant profiles and, consistently with the notation above, let this player be player 2. That is, consider first the profiles

$\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ and $\rho = (\rho_1, \rho_2)$ such that $\hat{\rho}_2 > \rho_2 \geq \hat{\rho}_1 > \rho_1$. Then

$$\begin{aligned}\Gamma_1^{UEC}(\rho, \hat{\rho}) &\equiv \left[V_1(\hat{\rho}_1; (\hat{\rho}_1, \rho_2)) - V_1(\rho_1; (\hat{\rho}_1, \rho_2)) \right] - \left[V_1(\hat{\rho}_1; (\rho_1, \rho_2)) - V_1(\rho_1; (\rho_1, \rho_2)) \right] \\ &= \frac{\sum_{k=\rho_1}^{\hat{\rho}_1} \sigma_k^2}{(1-\beta)^2} - (1+\beta)^2 \sum_{k=\rho_1}^{\hat{\rho}_1} \sigma_k^2 = \frac{1-(1-\beta^2)^2}{(1-\beta)^2} \sum_{k=\rho_1}^{\hat{\rho}_1} \sigma_k^2 > 0.\end{aligned}$$

Next, note that

$$\Gamma_2^{UEC}(\rho, \hat{\rho}) = \left[V_2(\hat{\rho}_2; (\rho_1, \hat{\rho}_2)) - V_2(\rho_2; (\rho_1, \hat{\rho}_2)) \right] - \left[V_2(\hat{\rho}_2; (\rho_1, \rho_2)) - V_2(\rho_2; (\rho_1, \rho_2)) \right] = 0.$$

Hence UEC holds strictly for player 1 and weakly for player 2.

Now consider increasing differences. Using the results above, we have that

$$\Gamma_1^{ID}(\rho, \hat{\rho}) = \left[V_1(\hat{\rho}_1; (\hat{\rho}_1, \hat{\rho}_2)) - V_1(\rho_1; (\hat{\rho}_1, \hat{\rho}_2)) \right] - \left[V_1(\hat{\rho}_1; (\hat{\rho}_1, \rho_2)) - V_1(\rho_1; (\hat{\rho}_1, \rho_2)) \right] = 0$$

and

$$\Gamma_2^{ID}(\rho, \hat{\rho}) = \left[V_2(\hat{\rho}_2; (\hat{\rho}_1, \hat{\rho}_2)) - V_2(\rho_2; (\hat{\rho}_1, \hat{\rho}_2)) \right] - \left[V_2(\hat{\rho}_2; (\rho_1, \hat{\rho}_2)) - V_2(\rho_2; (\rho_1, \hat{\rho}_2)) \right] = 0.$$

Combining the results for UEC with those for ID, we conclude that EC holds strictly for player 1 but only weakly (as an equality) for player 2.

Next, consider profiles $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ and $\rho = (\rho_1, \rho_2)$ such that $\hat{\rho}_2 = \hat{\rho}_1 > \rho_2 = \rho_1$. That UEC holds as an equality for these profiles follows directly from the fact that the action of the player who is behind is invariant in the number of additional dimensions $\hat{\rho}_j - \rho_j$ that the player who is ahead is expected to explore. That $\Gamma_i^{ID}(\rho, \hat{\rho}) > 0$ if $\beta > 0$ and $\Gamma_i^{ID}(\rho, \hat{\rho}) < 0$ if $\beta < 0$ follows from the fact that the marginal gross value of learning any dimension $k > \rho_i$ is equal to $\sigma_k^2 / (1-\beta)^2$ when learnt jointly and σ_k^2 when learnt alone. Q.E.D.

Proof of Proposition 5. First, we prove that all equilibria are symmetric. To see this, use the derivations in the proof of Proposition 4 to observe that, in any equilibrium in which cognition is symmetric, so are the stage-2 equilibrium actions. Hence, if an asymmetric equilibrium exists, it must feature asymmetric cognition choices. Thus assume an equilibrium exists in which $\rho_1 < \rho_2$. The arguments in the proof of Proposition 4 imply that, in any such equilibrium, the two players' ex-ante expected payoffs (disregarding again all external effects that have no influence on individual best responses) are given by

$$V_1(\rho_1; \rho) - C(\rho_1) = \frac{1 + \sum_{k=1}^{\rho_1} \sigma_k^2}{(1-\beta)^2} - \sum_{k=1}^{\rho_1} c_k$$

and

$$V_2(\rho_2; \rho) - C(\rho_2) = V_1(\rho_1; \rho) - C(\rho_1) + \sum_{k=\rho_1+1}^{\rho_2} (\sigma_k^2 - c_k).$$

Now suppose that player 1 deviates to cognition $\rho_1 < \rho'_1 \leq \rho_2$. The derivations in the proof of Proposition 4 imply that her payoff is equal to

$$V_1(\rho'_1; \rho) - C(\rho'_1) = V_1(\rho_1; \rho) - C(\rho_1) + \sum_{k=\rho_1+1}^{\rho'_1} ((1+\beta)^2 \sigma_k^2 - c_k).$$

The optimality of player 2's behavior in the putative equilibrium implies that $\sum_{k=\rho_1+1}^{\rho_1'} (\sigma_k^2 - c_k) > 0$. Because $\beta > 0$, we then have that player 1 has a profitable deviation. Hence, all equilibria are symmetric.

Next, observe that, in any (symmetric) equilibrium with cognitive profile ρ in which the depth of reasoning is $\rho_1 = \rho_2 = k^*$, the players' net equilibrium payoffs are equal to

$$V_i(k^*; (k^*, k^*)) - C(k^*) = \frac{1 + \sum_{k=1}^{k^*} \sigma_k^2}{(1 - \beta)^2} - \sum_{k=1}^{k^*} c_k,$$

$i = 1, 2$. For the two players not to have profitable deviations, it must be that neither a unilateral increase nor a unilateral decrease in cognition increases the players' payoff. Suppose that player 2 deviates to cognition $\rho_2 \geq k^* + 1$. The derivations in the proof of Proposition 4 imply that her net payoff following the deviation is equal to

$$V_i(k^*; (k^*, k^*)) - C(k^*) + \sum_{k=k^*+1}^{\rho_2} [\sigma_k^2 - c_k]$$

where the term in square brackets describes the net gain/loss. Because σ_k^2/c_k is decreasing, such deviations are unprofitable, for all $\rho_2 > k^*$, if and only if $\sigma_{k^*+1}^2 - c_{k^*+1} \leq 0$, i.e., if and only if $k^* \geq \underline{k} \equiv \min \left\{ k \mid \sigma_k^2 \leq c_k \right\}$.

Next, suppose that player 1 deviates to cognition $\rho_1' \leq k^* - 1$. The derivations in the proof of Proposition 4 imply that her expected net payoff following the deviation is equal to

$$V_i(k^* - 1; (k^* - 1, k^* - 1)) - C(k^* - 1) = V_i(k^*; (k^*, k^*)) - C(k^*) - \sum_{k=\rho_1'+1}^{k^*} \left[\frac{\sigma_k^2}{(1 - \beta)^2} - c_k \right].$$

where the term in brackets describes the net gain/loss. Again, because σ_k^2/c_k is decreasing, such deviations are unprofitable, no matter ρ_1' , if and only if $\frac{\sigma_{k^*}^2}{(1 - \beta)^2} \geq c_{k^*}$, which is equivalent to

$$k^* \leq \bar{k}(\beta) \equiv \max \left\{ k \mid \frac{\sigma_k^2}{(1 - \beta)^2} \geq c_k \right\}.$$

Hence, $k^* \in [\underline{k}, \bar{k}(\beta)]$ is both necessary and sufficient for a symmetric equilibrium with depth of cognition k^* .

Now assume that there are no external payoff effects (as explained in the main text, this is the case when the function ψ takes the form $\psi(a_{-i}, \omega) = h(\omega) + \beta \bar{a}_{-i}^2$). That equilibria are Pareto ranked then follows from the fact that, in the symmetric equilibrium with depth of knowledge k^* , the equilibrium payoffs (net of the term $\mathbb{E}[h(\omega)]$ which is exogenous to the agents' behavior) are given by

$$V_i(k^*; (k^*, k^*)) - C(k^*) = \frac{1 + \sum_{k=1}^{k^*} \sigma_k^2}{(1 - \beta)^2} - \sum_{k=1}^{k^*} c_k$$

which are increasing in k^* .

Finally, that any symmetric equilibrium with depth of knowledge $k^* \in [\underline{k}, \bar{k}(\beta)]$ is also an equilibrium in the game in which the players can select which dimensions to pay attention to follows from the fact that σ_k^2/c_k is decreasing in k . Q.E.D.

Proof of Proposition 6. From the arguments in the proof of Proposition 5 observe that, if a symmetric equilibrium exists, its cognitive level k^* must satisfy

$$\sigma_{k^*+1}^2 - c_{k^*+1} \leq 0 \leq \frac{\sigma_{k^*}^2}{(1-\beta)^2} - c_{k^*},$$

or, equivalently,

$$\frac{\sigma_{k^*+1}^2}{c_{k^*+1}} \leq 1 \leq \frac{\sigma_{k^*}^2}{(1-\beta)^2 c_{k^*}}. \quad (20)$$

The first inequality guarantees that a player does not benefit from learning the $k^* + 1$ dimension, whereas the second inequality that he does not benefit from limiting her knowledge to the $k^* - 1$ dimension, when both players are expected to learn k^* dimensions. Because σ_k^2/c_k is strictly decreasing, there can be at most one k^* satisfying the above two inequalities. Furthermore, when such a k^* exists, existence of a symmetric equilibrium also follows from the fact that σ_k^2/c_k is strictly decreasing, which implies that, if no local deviations are profitable, then nor are any of the global ones (the arguments are the same as in the proof of Proposition 5).

Next, consider asymmetric equilibria. Suppose the two players are expected to select $\rho = (\rho_1, \rho_2)$, with $\rho_1 < \rho_2$. Then player 2's payoff from choosing any cognition level $\rho'_2 \geq \rho_1$ is equal to

$$V_1(\rho_1; \rho) + \sum_{k=\rho_1+1}^{\rho'_2} \sigma_k^2 - \sum_{k=1}^{\rho'_2} c_k.$$

Hence, for player 2 to choose $\rho_2 > \rho_1$, it must be that

$$\sigma_{\rho_2+1}^2 - c_{\rho_2+1} \leq 0 \leq \sigma_{\rho_2}^2 - c_{\rho_2},$$

or, equivalently,

$$\frac{\sigma_{\rho_2+1}^2}{c_{\rho_2+1}} \leq 1 \leq \frac{\sigma_{\rho_2}^2}{c_{\rho_2}}. \quad (21)$$

The above double inequality, along with the monotonicity of σ_k^2/c_k in k , implies that, when player 1 chooses cognition ρ_1 , player 2 prefers cognition ρ_2 to any cognition $\rho'_2 \geq \rho_1$. Below we identify conditions under which, when the two players are expected to select cognition $\rho = (\rho_1, \rho_2)$, with $\rho_1 < \rho_2$, player 1 prefers cognition ρ_1 to cognition $\rho'_1 < \rho_1$. As shown in the proof of Proposition 4, the same conditions imply that player 2 also prefers ρ_1 to any $\rho'_2 < \rho_1$, and hence, by transitivity, ρ_2 to any $\rho'_2 < \rho_1$.³⁶

Next, consider player 1 (the one who selects the lowest cognitive level). As shown in the proof of Proposition 4, when the two players are expected to select cognition $\rho = (\rho_1, \rho_2)$, with $\rho_1 < \rho_2$, if player 1 were to increase her cognition to ρ'_1 , with $\rho_1 < \rho'_1 \leq \rho_2$, her payoff would be equal to

$$V_1(\rho_1; \rho) + (1+\beta)^2 \sum_{k=\rho_1+1}^{\rho'_1} \sigma_k^2 - \sum_{k=1}^{\rho'_1} c_k.$$

If, instead, she were to select cognition $\rho'_1 > \rho_2$, her payoff would be equal to

$$V_1(\rho_1; \rho) + (1+\beta)^2 \sum_{k=\rho_1+1}^{\rho_2} \sigma_k^2 + \sum_{k=\rho_2+1}^{\rho'_1} \sigma_k^2 - \sum_{k=1}^{\rho'_1} c_k.$$

³⁶Indeed, from the proof of Proposition 4, we have that the gross payoff that player 2 obtains from choosing cognition $\rho'_2 < \rho_1$ is equal to $\left[1 + \sum_{k=1}^{\rho'_2} \sigma_k^2\right] / (1-\beta)^2$ which is the same payoff that player 1 obtains from choosing cognition $\rho'_1 = \rho'_2$ when player 2 is expected to choose any cognitive level $\rho_2 \geq \rho_1$.

Finally, if she were to reduce her cognition to $\rho'_1 < \rho_1$, her payoff would be equal to

$$\frac{1 + \sum_{k=1}^{\rho'_1} \sigma_k^2}{(1 - \beta)^2} - \sum_{k=1}^{\rho'_1} c_k = V_1(\rho_1; \rho) - \sum_{k=1}^{\rho_1} c_k - \sum_{k=\rho'_1+1}^{\rho_1} \left(\frac{\sigma_k^2}{(1 - \beta)^2} - c_k \right).$$

Hence, for player 1 not to have profitable deviations, it must be that

$$(1 + \beta)^2 \sigma_{\rho_1+1}^2 - c_{\rho_1+1} \leq 0 \leq \frac{\sigma_{\rho_1}^2}{(1 - \beta)^2} - c_{\rho_1}$$

or, equivalently,

$$\frac{\sigma_{\rho_1+1}^2 (1 + \beta)^2}{c_{\rho_1+1}} \leq 1 \leq \frac{\sigma_{\rho_1}^2}{(1 - \beta)^2 c_{\rho_1}}. \quad (22)$$

The two inequalities in (22), along with the monotonicity of σ_k^2/c_k in k , also imply that player 1 prefers ρ_1 to any $\rho'_1 \leq \rho_2$. Furthermore, when paired with Condition (21), Condition (22) also implies that player 1 prefers ρ_1 to any $\rho'_1 > \rho_2$.

We thus conclude that Conditions (21) and (22) are both necessary and sufficient for any asymmetric equilibrium with cognition $\rho = (\rho_1, \rho_2)$, with $\rho_1 < \rho_2$. That, in any asymmetric equilibrium, the cognitive levels ρ_1 and ρ_2 belong to $[\underline{k}(\beta), \bar{k}]$ then follows from the above results along with the fact that $\rho_1 < \rho_2$.

Finally, consider the equilibrium payoffs. In any asymmetric equilibrium, player 1's equilibrium payoff is equal to

$$\frac{1 + \sum_{k=1}^{\rho_1} \sigma_k^2}{(1 - \beta)^2} - \sum_{k=1}^{\rho_1} c_k,$$

whereas player 2's equilibrium payoff is equal to

$$\frac{1 + \sum_{k=1}^{\rho_1} \sigma_k^2}{(1 - \beta)^2} - \sum_{k=1}^{\rho_1} c_k + \sum_{k=\rho_1+1}^{\rho_2} (\rho_k^2 - c_k).$$

It is then easy to see that player 1's equilibrium payoff is increasing in her depth of knowledge ρ_1 , whereas player 2's equilibrium payoff is decreasing in player 1's depth of knowledge and, in case there are multiple solutions ρ_2 to the double inequality in (21) is invariant in ρ_2 . It is also easy to see that the sum of the two players' payoff is maximal under the equilibrium featuring the lowest cognition for player 1 (the one learning the smaller number of dimensions). Q.E.D.

Proof of Proposition 7. Suppose that, in the stage-2 game, player i expects player j to follow a strategy that, for any $s_j = (x_j, y)$, selects the action

$$a_j(s_j) = m_j x_j + k_j y$$

with probability one, for some scalars m_j and k_j . Given her cognitive choice ρ_i , for any $s_i = (x_i, y)$, player i 's best response then consists in selecting the action

$$\begin{aligned} a_i &= (1 - \beta) \mathbb{E}[\omega | s_i; \rho_i] + \beta \mathbb{E}[a_j(s_j) | s_i; \rho_i] \\ &= (1 - \beta) \left[\frac{\rho_i}{\rho_i + h} x_i + \frac{h}{\rho_i + h} y \right] + \beta m_j \left[\frac{\rho_i}{\rho_i + h} x_i + \frac{h}{\rho_i + y} y \right] + \beta k_j y \end{aligned}$$

with probability one. Player i 's optimal action is thus also linear in x_i and y , with coefficients m_i and k_i given by

$$m_i = \frac{(1-\beta)\rho_i}{\rho_i+h} + \frac{\beta\rho_i}{\rho_i+h}m_j$$

and

$$k_i = \frac{(1-\beta)h + \beta hm_j}{\rho_i+h} + \beta k_j$$

Also note that, when $k_j = 1 - m_j$, then $k_i = 1 - m_i$. Iterating over the two players' best responses to find the fixed point, we then have that, for any cognitive profile $\rho = (\rho_i, \rho_j)$, there exists a unique linear continuation equilibrium for the stage-2 game and is such that, for any $s_i = (x_i, y)$, $i = 1, 2$, player i selects with probability one the action $a_i^\rho(s_i) = m_i^\rho x_i + k_i^\rho y$, where³⁷

$$m_i^\rho = (1-\beta)\rho_i \frac{\rho_j(1+\beta) + h}{(\rho_i+h)(\rho_j+h) - \beta^2\rho_i\rho_j}$$

and

$$k_i^\rho = h \left\{ \frac{\rho_j + h + \beta\rho_i}{(\rho_i+h)(\rho_j+h) - \beta^2\rho_i\rho_j} \right\} = 1 - m_i^\rho.$$

Expectation Conformity.

Next observe that, for any $(\omega, \varepsilon, \eta_1, \eta_2)$, and any (m_1, m_2) , when, in the stage-2 game, for any $s_i = (x_i, y)$, each player $i = 1, 2$ selects with probability one the action $a_i = m_i x_i + (1 - m_i)y$, then

$$a_i - \omega = m_i \eta_i + (1 - m_i)\varepsilon$$

and

$$a_i - a_j = m_i \eta_i - m_j \eta_j - (m_i - m_j)\varepsilon.$$

This implies that, when the two players are expected to select the cognitive profile $\rho = (\rho_i, \rho_j)$ and, instead, player i selects cognition ρ'_i , player i 's ex-ante expected (gross) payoff (disregarding the externality term ψ in (7)) is then equal to

$$\begin{aligned} V_i(\rho'_i; \rho) &= -(1-\beta) \left(m_i^{\rho'_i; \rho} \right)^2 \frac{1}{\rho'_i} - (1-\beta) \left(1 - m_i^{\rho'_i; \rho} \right)^2 \frac{1}{h} \\ &\quad - \beta \left(m_i^{\rho'_i; \rho} \right)^2 \frac{1}{\rho'_i} - \beta \left(m_j^\rho \right)^2 \frac{1}{\rho_j} - \beta \left(m_i^{\rho'_i; \rho} - m_j^\rho \right)^2 \frac{1}{h} \end{aligned}$$

where

$$m_j^\rho = \frac{(1-\beta)\rho_j(\rho_i+h) + \beta\rho_i\rho_j(1-\beta)}{(\rho_i+h)(\rho_j+h) - \beta^2\rho_i\rho_j}$$

is the sensitivity of player j 's stage-2 action to her private information when the two players are expected to select the cognitive profile $\rho = (\rho_i, \rho_j)$ and where

$$m_i^{\rho'_i; \rho} = \frac{(1-\beta)\rho'_i}{\rho'_i+h} + \frac{\beta\rho'_i}{\rho'_i+h}m_j^\rho$$

³⁷The formulas for the equilibrium sensitivities are obtained after various algebraic simplifications which are omitted for brevity.

is the sensitivity of player i 's stage-2 action to his private information when the two players are expected to select the cognitive profile $\rho = (\rho_i, \rho_j)$ and, instead, player i selects cognition ρ'_i .

Simplifying, then note that $V_i(\rho'_i; \rho)$ can be rewritten as

$$V_i(\rho'_i; \rho) = - \left(m_i^{\rho'_i; \rho} \right)^2 \frac{1}{\rho'_i} - \beta \left(m_j^\rho \right)^2 \frac{1}{\rho_j} - \left[1 - \beta + \left(m_i^{\rho'_i; \rho} \right)^2 - 2(1 - \beta) m_i^{\rho'_i; \rho} + \beta \left(m_j^\rho \right)^2 - 2\beta m_i^{\rho'_i; \rho} m_j^\rho \right] \frac{1}{h}$$

Now to see whether this game satisfies UEC, ID, and EC, take any pair of cognitive levels for player i , $\hat{\rho}_i$ and ρ_i , and let $\rho' = (\rho'_i, \rho'_j)$ and $\rho'' = (\rho''_i, \rho''_j)$ be two arbitrary cognitive profiles.

Then let

$$D \equiv \left[V_i(\hat{\rho}_i; \rho') - V_i(\rho_i; \rho') \right] - \left[V_i(\hat{\rho}_i; \rho'') - V_i(\rho_i; \rho'') \right].$$

Observe that UEC holds if $D \geq 0$ for $\rho' = (\hat{\rho}_i, \rho_j)$ and $\rho'' = (\rho_i, \rho_j)$, ID holds if $D \geq 0$ for $\rho' = (\hat{\rho}_i, \hat{\rho}_j)$ and $\rho'' = (\hat{\rho}_i, \rho_j)$, and EC holds if $D \geq 0$ for $\rho' = (\hat{\rho}_i, \hat{\rho}_j)$ and $\rho'' = (\rho_i, \rho_j)$.

Using the characterization of the V_i functions above, we have that³⁸

$$\begin{aligned} D &= \left[\left(m_i^{\rho'_i; \rho'} \right)^2 - \left(m_i^{\rho'_i; \rho''} \right)^2 \right] \left(\frac{1}{\rho_i} + \frac{1}{h} \right) - \left[\left(m_i^{\hat{\rho}_i; \rho'} \right)^2 - \left(m_i^{\hat{\rho}_i; \rho''} \right)^2 \right] \left(\frac{1}{\hat{\rho}_i} + \frac{1}{h} \right) \\ &+ 2 \left(m_i^{\hat{\rho}_i; \rho'} - m_i^{\rho'_i; \rho'} \right) \left[1 - \beta + \beta m_j^{\rho'} \right] \frac{1}{h} - 2 \left(m_i^{\hat{\rho}_i; \rho''} - m_i^{\rho'_i; \rho''} \right) \left[1 - \beta + \beta m_j^{\rho''} \right] \frac{1}{h}. \end{aligned}$$

Next, use the structure of the best responses derived above, to observe that

$$m_i^{\hat{\rho}_i; \rho'} - m_i^{\rho'_i; \rho'} = h \frac{(\hat{\rho}_i - \rho_i) \left(1 - \beta + \beta m_j^{\rho'} \right)}{(\hat{\rho}_i + h)(\rho_i + h)}$$

and

$$m_i^{\hat{\rho}_i; \rho''} - m_i^{\rho'_i; \rho''} = h \frac{(\hat{\rho}_i - \rho_i) \left(1 - \beta + \beta m_j^{\rho''} \right)}{(\hat{\rho}_i + h)(\rho_i + h)}.$$

Replacing the above expressions into the formulas for D , and simplifying, we then have that

$$\begin{aligned} D &= \left[\left(m_i^{\rho'_i; \rho'} \right)^2 - \left(m_i^{\rho'_i; \rho''} \right)^2 \right] \left(\frac{1}{\rho_i} + \frac{1}{h} \right) - \left[\left(m_i^{\hat{\rho}_i; \rho'} \right)^2 - \left(m_i^{\hat{\rho}_i; \rho''} \right)^2 \right] \left(\frac{1}{\hat{\rho}_i} + \frac{1}{h} \right) \\ &+ \frac{2(\hat{\rho}_i - \rho_i) \left[\left(1 - \beta + \beta m_j^{\rho'} \right)^2 - \left(1 - \beta + \beta m_j^{\rho''} \right)^2 \right]}{(\hat{\rho}_i + h)(\rho_i + h)}. \end{aligned}$$

Using the fact that

$$m_i^{\rho'_i; \rho'} - m_i^{\rho'_i; \rho''} = \frac{\beta \rho_i}{\rho_i + h} \left(m_j^{\rho'} - m_j^{\rho''} \right)$$

and

$$m_i^{\hat{\rho}_i; \rho'} - m_i^{\hat{\rho}_i; \rho''} = \frac{\beta \hat{\rho}_i}{\hat{\rho}_i + h} \left(m_j^{\rho'} - m_j^{\rho''} \right)$$

³⁸Again, the formula for D is obtained after various algebraic simplifications that are omitted here for brevity.

after a few simplification, we then have that

$$D = \frac{\beta}{h} \left(m_j^{\rho'} - m_j^{\rho''} \right) \left(m_i^{\rho_i; \rho'} + m_i^{\rho_i; \rho''} - m_i^{\hat{\rho}_i; \rho'} - m_i^{\hat{\rho}_i; \rho''} \right) \\ + \frac{2(\hat{\rho}_i - \rho_i) \left[(1 - \beta + \beta m_j^{\rho'})^2 - (1 - \beta + \beta m_j^{\rho''})^2 \right]}{(\hat{\rho}_i + h)(\rho_i + h)}.$$

Next observe that

$$m_i^{\rho_i; \rho'} + m_i^{\rho_i; \rho''} = \frac{2(1 - \beta)\rho_i}{\rho_i + h} + \frac{\beta\rho_i}{\rho_i + h} \left(m_j^{\rho'} + m_j^{\rho''} \right)$$

and

$$m_i^{\hat{\rho}_i; \rho'} + m_i^{\hat{\rho}_i; \rho''} = \frac{2(1 - \beta)\hat{\rho}_i}{\hat{\rho}_i + h} + \frac{\beta\hat{\rho}_i}{\hat{\rho}_i + h} \left(m_j^{\rho'} + m_j^{\rho''} \right)$$

from which we obtain that

$$m_i^{\rho_i; \rho'} + m_i^{\rho_i; \rho''} - m_i^{\hat{\rho}_i; \rho'} - m_i^{\hat{\rho}_i; \rho''} = \\ = \frac{2(1 - \beta)\rho_i(\hat{\rho}_i + h) + \beta\rho_i(\hat{\rho}_i + h)(m_j^{\rho'} + m_j^{\rho''}) - 2(1 - \beta)\hat{\rho}_i(\rho_i + h) - \beta\hat{\rho}_i(\rho_i + h)(m_j^{\rho'} + m_j^{\rho''})}{(\rho_i + h)(\hat{\rho}_i + h)}.$$

The numerator is equal to

$$N = 2(1 - \beta)\rho_i h + \beta\rho_i h \left(m_j^{\rho'} + m_j^{\rho''} \right) - 2(1 - \beta)\hat{\rho}_i h - \beta\hat{\rho}_i h \left(m_j^{\rho'} + m_j^{\rho''} \right).$$

Hence,

$$m_i^{\rho_i; \rho'} + m_i^{\rho_i; \rho''} - m_i^{\hat{\rho}_i; \rho'} - m_i^{\hat{\rho}_i; \rho''} = -h \frac{2(1 - \beta)(\hat{\rho}_i - \rho_i) + \beta(\hat{\rho}_i - \rho_i) \left(m_j^{\rho'} + m_j^{\rho''} \right)}{(\rho_i + h)(\hat{\rho}_i + h)}$$

Replacing the above into the formula for D , we have that

$$D = \frac{\hat{\rho}_i - \rho_i}{(\rho_i + h)(\hat{\rho}_i + h)} \beta \left(m_j^{\rho'} - m_j^{\rho''} \right) \left[2(1 - \beta) + \beta \left(m_j^{\rho'} + m_j^{\rho''} \right) \right]$$

from which we obtain that

$$D \stackrel{\text{sgn}}{=} \beta(\hat{\rho}_i - \rho_i) \left(m_j^{\rho'} - m_j^{\rho''} \right).$$

We are now ready to establish the properties in the proposition.

To see that this game satisfies UEC, take $\rho' = (\hat{\rho}_i, \rho_j)$ and $\rho'' = (\rho_i, \rho_j)$. Using the fact that, for any $\rho = (\rho_i, \rho_j)$,

$$m_j^\rho = \frac{(1 - \beta)\rho_j(\rho_i + h) + \beta\rho_i\rho_j(1 - \beta)}{(\rho_i + h)(\rho_j + h) - \beta^2\rho_i\rho_j}$$

we have that

$$\frac{\partial m_j^\rho}{\partial \rho_i} \stackrel{\text{sgn}}{=} \beta.$$

We thus conclude that, independently of the sign of β , $D \geq 0$, which implies that the game satisfies UEC.

Next, to see whether this game satisfies ID, take $\rho' = (\hat{\rho}_i, \hat{\rho}_j)$ and $\rho'' = (\hat{\rho}_i, \rho_j)$. Using the fact that, for any $\rho = (\rho_i, \rho_j)$, and any $\beta \in (-1, +1)$, m_j^ρ is non-decreasing in ρ_j , we then have that

$$D \stackrel{\text{sgn}}{=} \beta(\hat{\rho}_i - \rho_i)(\hat{\rho}_j - \rho_j).$$

We thus have that, for any pair of information structures $\hat{\rho} = (\hat{\rho}_i, \hat{\rho}_j)$ and $\rho = (\rho_i, \rho_j)$, this game exhibits ID if and only if $\beta(\hat{\rho}_i - \rho_i)(\hat{\rho}_j - \rho_j) \geq 0$.

Finally, to see whether this game satisfies EC, take $\rho' = (\hat{\rho}_i, \hat{\rho}_j)$ and $\rho'' = (\rho_i, \rho_j)$. We then have that

$$\frac{m_j^{\rho'} - m_j^{\rho''}}{1 - \beta} = \frac{\hat{\rho}_j(\hat{\rho}_i + h) + \beta\hat{\rho}_i\hat{\rho}_j}{(\hat{\rho}_i + h)(\hat{\rho}_j + h) - \beta^2\hat{\rho}_i\hat{\rho}_j} - \frac{\rho_j(\rho_i + h) + \beta\rho_i\rho_j}{(\rho_i + h)(\rho_j + h) - \beta^2\rho_i\rho_j}$$

from which we obtain that

$$\frac{m_j^{\rho'} - m_j^{\rho''}}{1 - \beta} = \frac{[\hat{\rho}_j(\hat{\rho}_i + h) + \beta\hat{\rho}_i\hat{\rho}_j][(\rho_i + h)(\rho_j + h) - \beta^2\rho_i\rho_j] - [\rho_j(\rho_i + h) + \beta\rho_i\rho_j][(\hat{\rho}_i + h)(\hat{\rho}_j + h) - \beta^2\hat{\rho}_i\hat{\rho}_j]}{[(\hat{\rho}_i + h)(\hat{\rho}_j + h) - \beta^2\hat{\rho}_i\hat{\rho}_j][(\rho_i + h)(\rho_j + h) - \beta^2\rho_i\rho_j]}.$$

The denominator is always positive, so focus on the numerator. This is equal to

$$\begin{aligned} n &= [\hat{\rho}_j(\hat{\rho}_i + h) + \beta\hat{\rho}_i\hat{\rho}_j][(\rho_i + h)(\rho_j + h) - \beta^2\rho_i\rho_j] - [\rho_j(\rho_i + h) + \beta\rho_i\rho_j][(\hat{\rho}_i + h)(\hat{\rho}_j + h) - \beta^2\hat{\rho}_i\hat{\rho}_j] \\ &= [\hat{\rho}_j(\hat{\rho}_i + h) + \beta\hat{\rho}_i\hat{\rho}_j][(\rho_i + h)(\rho_j + h)] - \beta^2\rho_i\rho_j\hat{\rho}_j h \\ &\quad - [\rho_j(\rho_i + h) + \beta\rho_i\rho_j][(\hat{\rho}_i + h)(\hat{\rho}_j + h)] + \beta^2\hat{\rho}_i\hat{\rho}_j\rho_j h. \end{aligned}$$

Simplifying, we have that

$$\begin{aligned} \frac{n}{h} &= (\hat{\rho}_j - \rho_j)[\hat{\rho}_i\rho_i(1 + \beta) + (\hat{\rho}_i + \rho_i)h] + (\hat{\rho}_i - \rho_i)\hat{\rho}_j\rho_j(\beta + \beta^2) \\ &\quad + \hat{\rho}_i\hat{\rho}_j\beta h + h^2\hat{\rho}_j - \rho_i\rho_j\beta h - h^2\rho_j \end{aligned}$$

Hence, EC holds with respect to $\hat{\rho} = (\hat{\rho}_i, \hat{\rho}_j)$ and $\rho = (\rho_i, \rho_j)$ if and only if

$$\begin{aligned} &\beta(\hat{\rho}_j - \rho_j)(\hat{\rho}_i - \rho_i)[\hat{\rho}_i\rho_i(1 + \beta) + (\hat{\rho}_i + \rho_i)h] + \beta^2(\hat{\rho}_i - \rho_i)^2\hat{\rho}_j\rho_j(1 + \beta) \\ &\quad + \beta h[\beta(\hat{\rho}_i\hat{\rho}_j - \rho_i\rho_j) + h(\hat{\rho}_j - \rho_j)](\hat{\rho}_i - \rho_i) \geq 0 \end{aligned}$$

which, after some simplifications, can be rewritten as

$$\beta(\hat{\rho}_j - \rho_j)(\hat{\rho}_i - \rho_i)(\hat{\rho}_i + h)[\rho_i(1 + \beta) + h] + \beta^2(\hat{\rho}_i - \rho_i)^2\hat{\rho}_j[\rho_j(1 + \beta) + h] \geq 0 .$$

Q.E.D.

Proof of Proposition 8. First, observe that, given $s_i = (s_i^P, s_i^S)$,

$$\mathbb{E}[\omega|s_i] = \frac{1}{1 + h_i^S} s_i^P + \frac{h_i^S}{1 + h_i^S} s_i^S$$

where 1 is the precision of agent i 's primary signal and where

$$h_i^S = [\text{var}(\varepsilon_j + \eta_i)]^{-1} = \frac{\rho_i}{1 + \rho_i}$$

is the precision of his secondary signal. Hence,

$$\mathbb{E}[\omega | s_i] = \frac{1 + \rho_i}{1 + 2\rho_i} s_i^P + \frac{\rho_i}{1 + 2\rho_i} s_i^S.$$

Similarly,

$$\mathbb{E}[s_j^S | s_i] = s_i^P.$$

Finally,

$$\begin{aligned} \mathbb{E}[s_j^P | s_i] &= \mathbb{E}\left[\mathbb{E}[s_j^P | s_i, \omega] | s_i\right] = \mathbb{E}\left[\omega + \mathbb{E}[\varepsilon_j | s_i, \omega] | s_i\right] = \mathbb{E}\left[\omega + \frac{\rho_i}{1 + \rho_i}(s_i^S - \omega) | s_i\right] \\ &= \frac{1}{1 + \rho_i} \mathbb{E}[\omega | s_i] + \frac{\rho_i}{1 + \rho_i} s_i^S = \frac{1}{1 + 2\rho_i} s_i^P + \frac{2\rho_i}{1 + 2\rho_i} s_i^S. \end{aligned}$$

Now, suppose that, in the stage-2 game player i expects player j to follow a strategy that, for any $s_j = (s_j^P, s_j^S)$ selects with probability one the action

$$a_j(s_j) = m_j s_j^P + k_j s_j^S$$

for some scalars m_j and k_j . Given ρ_i , for any $s_i = (s_i^P, s_i^S)$, player i 's best response in the stage-2 game then consists in choosing with certainty the action

$$\begin{aligned} a_i &= (1 - \beta) \mathbb{E}[\omega | s_i] + \beta \mathbb{E}[a_j(s_j) | s_i] \\ &= (1 - \beta) \frac{1 + \rho_i}{1 + 2\rho_i} s_i^P + (1 - \beta) \frac{\rho_i}{1 + 2\rho_i} s_i^S \\ &\quad + \beta m_j \left[\frac{1}{1 + 2\rho_i} s_i^P + \frac{2\rho_i}{1 + 2\rho_i} s_i^S \right] + \beta k_j s_i^P. \end{aligned}$$

Player i 's best optimal action $a_i(s_i) = m_i s_i^P + k_i s_i^S$ is thus also linear in s_i with coefficients m_i and k_i given by

$$m_i = \frac{(1 - \beta)(1 + \rho_i)}{1 + 2\rho_i} + \frac{\beta}{1 + 2\rho_i} m_j + \beta k_j$$

and

$$k_i = \frac{(1 - \beta)\rho_i}{1 + 2\rho_i} + \frac{2\beta\rho_i}{1 + 2\rho_i} m_j.$$

Furthermore, as anticipated in the main text, when $k_j = 1 - m_j$, then $k_i = 1 - m_i$. In this case, player i 's best response is given by the linear action whose coefficients are given by

$$m_i = \frac{1 + \rho_i(1 + \beta) - 2\beta\rho_i m_j}{1 + 2\rho_i}$$

and

$$k_i = \rho_i \frac{1 - \beta + 2\beta m_j}{1 + 2\rho_i} = 1 - m_i.$$

Fixing the two players' cognitive structures, $\rho = (\rho_1, \rho_2)$, after some algebra, one can then verify that the stage-2 equilibrium actions are given by

$$m_j = \frac{1 + 2\rho_i + \rho_j(1 - \beta) + 2\rho_i\rho_j(1 - \beta^2)}{1 + 2(\rho_i + \rho_j) + 4\rho_i\rho_j(1 - \beta^2)}$$

and

$$k_j = \rho_j(1 + \beta) \left\{ \frac{1 + 2\rho_i(1 - \beta)}{1 + 2(\rho_i + \rho_j) + 4\rho_i\rho_j(1 - \beta^2)} \right\} = 1 - m_j.$$

Hence, for any cognitive profile $\rho = (\rho_i, \rho_j)$, the stage-2 game admits a unique linear continuation equilibrium and in such equilibrium each player $i = 1, 2$, for any $s_i = (s_i^P, s_i^S)$, selects with certainty the action $a_i^\rho(s_i) = m_i^\rho s_i^P + k_i^\rho s_i^S$, with

$$m_i^\rho = \frac{1 + 2\rho_j + \rho_i(1 - \beta) + 2\rho_i\rho_j(1 - \beta^2)}{1 + 2(\rho_i + \rho_j) + 4\rho_i\rho_j(1 - \beta^2)} \quad (23)$$

and

$$k_i^\rho = \rho_i(1 + \beta) \left\{ \frac{1 + 2\rho_j(1 - \beta)}{1 + 2(\rho_i + \rho_j) + 4\rho_i\rho_j(1 - \beta^2)} \right\} = 1 - m_i^\rho.$$

Expectation Conformity.

Next observe that, for any $(\omega, \varepsilon_1, \varepsilon_2, \eta_1, \eta_2)$, and any (m_1, m_2) , when, in the stage-2 game, for any $s_i = (s_i^P, s_i^S)$, each player $i = 1, 2$ selects with probability one the action $a_i = m_i s_i^P + (1 - m_i) s_i^S$, then

$$a_i - \omega = m_i \varepsilon_i + (1 - m_i)(\varepsilon_j + \eta_i)$$

and

$$a_i - a_j = [m_i - (1 - m_j)] \varepsilon_i + [(1 - m_i) - m_j] \varepsilon_j + (1 - m_i) \eta_i - (1 - m_j) \eta_j.$$

In turn, this implies that, when the two players are expected to select the cognitive profile $\rho = (\rho_i, \rho_j)$ and, instead, player i selects cognition ρ'_i , player i 's ex-ante expected gross payoff (as usual, disregarding the externality term ψ in (7)) is equal to

$$\begin{aligned} V_i(\rho'_i; \rho) &= -(1 - \beta) \left(m_i^{\rho'_i; \rho} \right)^2 - (1 - \beta) \left(1 - m_i^{\rho'_i; \rho} \right)^2 \left(1 + \frac{1}{\rho'_i} \right) \\ &\quad - \beta \left[m_i^{\rho'_i; \rho} - (1 - m_j^\rho) \right]^2 - \beta \left[(1 - m_i^{\rho'_i; \rho}) - m_j^\rho \right]^2 - \beta \left(1 - m_i^{\rho'_i; \rho} \right)^2 \frac{1}{\rho'_i} - \beta \left(1 - m_j^\rho \right)^2 \frac{1}{\rho_j} \end{aligned}$$

where m_j^ρ is as in ((23)) but applied to player j , and where

$$m_i^{\rho'_i; \rho} = \frac{1 + \rho'_i(1 + \beta) - 2\beta\rho'_i m_j^\rho}{1 + 2\rho'_i} \quad (24)$$

is player i 's optimal sensitivity to his primary signal when the two players are expected to select the cognitive profile $\rho = (\rho_i, \rho_j)$ and, instead, player i selects cognition ρ'_i [as shown above, the sensitivities to the secondary signals are equal to $k_j^\rho = 1 - m_j^\rho$ and $k_i^{\rho'_i; \rho} = 1 - m_i^{\rho'_i; \rho}$, respectively].

To see whether this game satisfies unilateral expectation conformity (UEC) and increasing differences (ID), as in the proof of Proposition (7), take any pair of cognitive levels for player i , $\hat{\rho}_i$ and ρ_i , and let $\rho' = (\rho'_i, \rho'_j)$ and $\rho'' = (\rho''_i, \rho''_j)$ be two arbitrary cognitive profiles.

Then let

$$D \equiv \left[V_i(\hat{\rho}_i; \rho') - V_i(\rho_i; \rho') \right] - \left[V_i(\hat{\rho}_i; \rho'') - V_i(\rho_i; \rho'') \right]$$

denote the difference in payoff differentials across the two profiles. Using the Characterization of the V_i function above, we have that

$$\begin{aligned} D &= 2(1 - \beta) \left(m_i^{\rho_i; \rho'} - m_i^{\rho_i; \rho''} \right) - 2(1 - \beta) \left(m_i^{\hat{\rho}_i; \rho'} - m_i^{\hat{\rho}_i; \rho''} \right) \\ &\quad + 4\beta m_j^{\rho''} \left(m_i^{\hat{\rho}_i; \rho''} - m_i^{\rho_i; \rho''} \right) + 4\beta m_j^{\rho'} \left(m_i^{\rho_i; \rho'} - m_i^{\hat{\rho}_i; \rho'} \right) \\ &\quad + \left(m_i^{\rho_i; \rho'} - m_i^{\rho_i; \rho''} \right) \left(m_i^{\rho_i; \rho'} + m_i^{\rho_i; \rho''} - 2 \right) \left[\frac{1+2\rho_i}{\rho_i} \right] \\ &\quad - \left(m_i^{\hat{\rho}_i; \rho'} - m_i^{\hat{\rho}_i; \rho''} \right) \left(m_i^{\hat{\rho}_i; \rho'} + m_i^{\hat{\rho}_i; \rho''} - 2 \right) \left[\frac{1+2\hat{\rho}_i}{\hat{\rho}_i} \right]. \end{aligned}$$

Now, using ((24)), we have that

$$m_i^{\rho_i; \rho'} - m_i^{\rho_i; \rho''} = -\frac{2\beta\rho_i}{1+2\rho_i} (m_j^{\rho'} - m_j^{\rho''})$$

and

$$m_i^{\hat{\rho}_i; \rho'} - m_i^{\hat{\rho}_i; \rho''} = -\frac{2\beta\hat{\rho}_i}{1+2\hat{\rho}_i} (m_j^{\rho'} - m_j^{\rho''}).$$

Replacing the above expressions into the formula for D above, after some algebra, we have that

$$D = 4\beta (\hat{\rho}_i - \rho_i) \frac{m_j^{\rho'} - m_j^{\rho''}}{(1+2\rho_i)(1+2\hat{\rho}_i)} \left(\beta m_j^{\rho'} + \beta m_j^{\rho''} + 1 - \beta \right).$$

Observe that, independently of the sign of β , $\beta m_j^{\rho'} + \beta m_j^{\rho''} + 1 - \beta > 0$. Hence,

$$D \stackrel{\text{sgn}}{\equiv} \beta (\hat{\rho}_i - \rho_i) \left(m_j^{\rho'} - m_j^{\rho''} \right).$$

Next observe that

$$m_j^{\rho'} - m_j^{\rho''} = \frac{n}{d}$$

where

$$\begin{aligned} n &\equiv \left[1 + 2\rho'_i + \rho'_j(1 - \beta) + 2\rho'_i\rho'_j(1 - \beta^2) \right] \left[1 + 2(\rho''_i + \rho''_j) \right] + 4\rho''_i\rho''_j(1 - \beta^2) \left[1 + 2\rho'_i + \rho'_j(1 - \beta) \right] \\ &\quad - \left[1 + 2\rho''_i + \rho''_j(1 - \beta) + 2\rho''_i\rho''_j(1 - \beta^2) \right] \left[1 + 2(\rho'_i + \rho'_j) \right] - 4\rho'_i\rho'_j(1 - \beta^2) \left[1 + 2\rho''_i + \rho''_j(1 - \beta) \right] \end{aligned}$$

and

$$d \equiv \left[1 + 2(\rho'_i + \rho'_j) + 4\rho'_i\rho'_j(1 - \beta^2) \right] \left[1 + 2(\rho''_i + \rho''_j) + 4\rho''_i\rho''_j(1 - \beta^2) \right]$$

Clearly, $d > 0$, whereas

$$\frac{n}{1+\beta} = \rho_j'' - \rho_j' + 2\rho_i'\rho_j'' - 2\rho_i''\rho_j' + 2\rho_i''\rho_j''(1-\beta) \left[2\rho_i' + 1 - 2\rho_j'\beta \right] - 2\rho_i'\rho_j'(1-\beta) \left[2\rho_i'' + 1 - 2\rho_j''\beta \right].$$

Hence,

$$D \stackrel{\text{sgn}}{=} \beta(\hat{\rho}_i - \rho_i) \frac{n}{1+\beta}.$$

To see that this game satisfies UEC, take $\rho' = (\hat{\rho}_i, \rho_j)$ and $\rho'' = (\rho_i, \rho_j)$. We then have that

$$\frac{n}{1+\beta} = 2\rho_j\beta [1 + 2\rho_j(1-\beta)] (\hat{\rho}_i - \rho_i).$$

We thus conclude that, irrespective of the sign of β , $D \geq 0$, which implies that the game satisfies UEC.

Next, to see whether this game satisfies ID, take $\rho' = (\hat{\rho}_i, \hat{\rho}_j)$ and $\rho'' = (\hat{\rho}_i, \rho_j)$. We then have that

$$\frac{n}{1+\beta} = -(\hat{\rho}_j - \rho_j) [1 + 2\hat{\rho}_i + 4\hat{\rho}_i^2(1-\beta) + 2\hat{\rho}_i(1-\beta)]$$

and hence that $D \stackrel{\text{sgn}}{=} -\beta(\hat{\rho}_i - \rho_i) (\hat{\rho}_j - \rho_j)$. We conclude that the game satisfies ID if and only if $\beta(\hat{\rho}_i - \rho_i) (\hat{\rho}_j - \rho_j) \leq 0$. Q.E.D.

Proof of Proposition 9. First, consider UEC. Given any pair of cognitive profiles $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ and $\rho = (\rho_1, \rho_2)$, the value for player 2 (the receiver) of going from cognition ρ_2 to cognition $\hat{\rho}_2$ when player 1's cognition is ρ_1 is invariant in the level of cognition that player 1 expects player 2 to select. Hence, UEC trivially holds for player 2, albeit never strictly. Next, consider player 1 (the persuader). Using the characterization of the ex-ante expected gross payoffs in the paragraphs preceding the proposition, we have that

$$\begin{aligned} \Gamma_1^{UEC}(\rho, \hat{\rho}) &\equiv \left[V_1(\hat{\rho}_1; (\hat{\rho}_1, \rho_2)) - V_1(\rho_1; (\hat{\rho}_1, \rho_2)) \right] - \left[V_1(\hat{\rho}_1; (\rho_1, \rho_2)) - V_1(\rho_1; (\rho_1, \rho_2)) \right] \\ &= F(0) [(1 - r^-(\rho_2))\bar{\omega}(\hat{\rho}_1, \rho_2) + r^-(\rho_2)\omega^-] + (1 - F(0)) [(1 - r^+(\hat{\rho}_1, \rho_2))\bar{\omega}(\hat{\rho}_1, \rho_2) + r^+(\hat{\rho}_1, \rho_2)\omega^+] \\ &\quad - F(0) [(1 - r^-(\rho_2))\bar{\omega}(\hat{\rho}_1, \rho_2) + r^-(\rho_2)\omega^-] - (1 - F(0)) [(1 - r^+(\rho_1, \rho_2))\bar{\omega}(\hat{\rho}_1, \rho_2) + r^+(\rho_1, \rho_2)\omega^+] \\ &\quad - F(0) [(1 - r^-(\rho_2))\bar{\omega}(\rho_1, \rho_2) + r^-(\rho_2)\omega^-] - (1 - F(0)) [(1 - r^+(\hat{\rho}_1, \rho_2))\bar{\omega}(\rho_1, \rho_2) + r^+(\hat{\rho}_1, \rho_2)\omega^+] \\ &\quad + F(0) [(1 - r^-(\rho_2))\bar{\omega}(\rho_1, \rho_2) + r^-(\rho_2)\omega^-] + (1 - F(0)) [(1 - r^+(\rho_1, \rho_2))\bar{\omega}(\rho_1, \rho_2) + r^+(\rho_1, \rho_2)\omega^+]. \end{aligned}$$

After simplifying, we have that

$$\Gamma_1^{UEC}(\rho, \hat{\rho}) = [1 - F(0)] [r^+(\hat{\rho}_1, \rho_2) - r^+(\rho_1, \rho_2)] [\bar{\omega}(\rho_1, \rho_2) - \bar{\omega}(\hat{\rho}_1, \rho_2)] \geq 0$$

where the inequality follows from the fact that

$$\bar{\omega}(\rho) = \frac{(1 - r^-(\rho_2))F(0)\omega^- + (1 - r^+(\rho_1, \rho_2))(1 - F(0))\omega^+}{(1 - r^-(\rho_2))F(0) + (1 - r^+(\rho_1, \rho_2))(1 - F(0))}$$

is decreasing in ρ_1 . Hence, UEC holds strictly for player 1, for any pair of cognitive profiles $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ and $\rho = (\rho_1, \rho_2)$ such that $\hat{\rho}_1 \neq \rho_1$ (irrespective of the sign of $\hat{\rho}_1 - \rho_1$).

Next, consider ID. Note that

$$\begin{aligned} \Gamma_1^{ID}(\rho, \hat{\rho}) &\equiv \left[V_1(\hat{\rho}_1; \hat{\rho}_1, \hat{\rho}_2) - V_1(\rho_1; \hat{\rho}_1, \hat{\rho}_2) \right] - \left[V_1(\hat{\rho}_1; \hat{\rho}_1, \rho_2) - V_1(\rho_1; \hat{\rho}_1, \rho_2) \right] \\ &= F(0) [(1 - r^-(\hat{\rho}_2))\bar{\omega}(\hat{\rho}_1, \hat{\rho}_2) + r^-(\hat{\rho}_2)\omega^-] + (1 - F(0)) [(1 - r^+(\hat{\rho}_1, \hat{\rho}_2))\bar{\omega}(\hat{\rho}_1, \hat{\rho}_2) + r^+(\hat{\rho}_1, \hat{\rho}_2)\omega^+] \\ &\quad - F(0) [(1 - r^-(\hat{\rho}_2))\bar{\omega}(\hat{\rho}_1, \hat{\rho}_2) + r^-(\hat{\rho}_2)\omega^-] - (1 - F(0)) [(1 - r^+(\rho_1, \hat{\rho}_2))\bar{\omega}(\hat{\rho}_1, \hat{\rho}_2) + r^+(\rho_1, \hat{\rho}_2)\omega^+] \\ &\quad - F(0) [(1 - r^-(\rho_2))\bar{\omega}(\hat{\rho}_1, \rho_2) + r^-(\rho_2)\omega^-] - (1 - F(0)) [(1 - r^+(\hat{\rho}_1, \rho_2))\bar{\omega}(\hat{\rho}_1, \rho_2) + r^+(\hat{\rho}_1, \rho_2)\omega^+] \\ &\quad + F(0) [(1 - r^-(\rho_2))\bar{\omega}(\hat{\rho}_1, \rho_2) + r^-(\rho_2)\omega^-] + (1 - F(0)) [(1 - r^+(\rho_1, \rho_2))\bar{\omega}(\hat{\rho}_1, \rho_2) + r^+(\rho_1, \rho_2)\omega^+]. \end{aligned}$$

After simplifying, we have that

$$\begin{aligned} \frac{\Gamma_1^{ID}(\rho, \hat{\rho})}{(1-F(0))} &= [r^+(\hat{\rho}_1, \hat{\rho}_2) - r^+(\rho_1, \hat{\rho}_2)] [\omega^+ - \bar{\omega}(\hat{\rho}_1, \hat{\rho}_2)] \\ &\quad - [r^+(\hat{\rho}_1, \rho_2) - r^+(\rho_1, \rho_2)] [\omega^+ - \bar{\omega}(\hat{\rho}_1, \rho_2)]. \end{aligned}$$

Hence, ID holds for player 1 with respect to the cognitive profiles $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ and $\rho = (\rho_1, \rho_2)$ if and only if Condition (16) in the proposition holds. It is easy to see that Condition (16) is satisfied when properties (a)-(c) in the proposition hold.

Finally, note that

$$\begin{aligned} \frac{\partial \bar{\omega}(\rho)}{\partial \rho_2} \stackrel{sgn}{=} \left[-\frac{dr^-(\rho_2)}{d\rho_2} F(0)\omega^- - \frac{\partial r^+(\rho_1, \rho_2)}{\partial \rho_2} (1 - F(0))\omega^+ \right] [(1 - r^-(\rho_2))F(0) + (1 - r^+(\rho_1, \rho_2))(1 - F(0))] \\ - [(1 - r^-(\rho_2))F(0)\omega^- + (1 - r^+(\rho_1, \rho_2))(1 - F(0))\omega^+] \left[-\frac{dr^-(\rho_2)}{d\rho_2} F(0) - \frac{\partial r^+(\rho_1, \rho_2)}{\partial \rho_2} (1 - F(0)) \right]. \end{aligned}$$

After some algebra, we have that

$$\frac{\partial \bar{\omega}(\rho)}{\partial \rho_2} \stackrel{sgn}{=} \frac{dr^-(\rho_2)}{d\rho_2} - \frac{\partial r^+(\rho_1, \rho_2)}{\partial \rho_2} - r^+(\rho_1, \rho_2) \frac{dr^-(\rho_2)}{d\rho_2} + \frac{\partial r^+(\rho_1, \rho_2)}{\partial \rho_2} r^-(\rho_2) .$$

Therefore, $\bar{\omega}(\rho)$ is decreasing in ρ_2 for example when $\frac{dr^-(\rho_2)}{d\rho_2} = \frac{\partial r^+(\rho_1, \rho_2)}{\partial \rho_2}$ and $r^+(\rho_1, \rho_2) \geq r^-(\rho_2)$, whereas $\bar{\omega}(\rho)$ is increasing in ρ_2 when $\frac{dr^-(\rho_2)}{d\rho_2} > \frac{\partial r^+(\rho_1, \rho_2)}{\partial \rho_2}$ and $r^+(\rho_1, \rho_2) \cong r^-(\rho_2)$, as claimed in the paragraphs in the main text preceding Proposition 9. Q.E.D.

Proof of Proposition 10. First, consider UEC. Using the expressions for $V_1(\rho'_1; \rho)$ and $V_2(\rho'_2; \rho)$ is the main text, we have that

$$\begin{aligned} \Gamma_1^{UEC}(\rho, \hat{\rho}) &= \frac{M(\hat{\rho}_1, \rho_2)h_\varepsilon}{M^2(\hat{\rho}_1, \rho_2)h_\varepsilon + R^2(\hat{\rho}_1, \rho_2)h_\omega} [M(\hat{\rho}_1, \rho_2)\omega_0 + A(\hat{\rho}_1, \rho_2) - M(\rho_1, \rho_2)\omega_0 - A(\rho_1, \rho_2)] \\ &\quad - \frac{M(\rho_1, \rho_2)h_\varepsilon}{M^2(\rho_1, \rho_2)h_\varepsilon + R^2(\rho_1, \rho_2)h_\omega} [M(\hat{\rho}_1, \rho_2)\omega_0 + A(\hat{\rho}_1, \rho_2) - M(\rho_1, \rho_2)\omega_0 - A(\rho_1, \rho_2)] \end{aligned}$$

and $\Gamma_2^{UEC}(\rho, \hat{\rho}) = 0$. Hence, for any pair of profiles $\rho = (\rho_1, \rho_2)$ and $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ with $\hat{\rho}_i \geq \rho_i$, $i = 1, 2$, we have that $\Gamma_1^{UEC}(\rho, \hat{\rho}) = 0$ if the function L is constant in ρ_1 . When, instead, L is strictly increasing in ρ_1 , then $\Gamma_1^{UEC}(\rho, \hat{\rho}) \geq 0$ if the function G is non-decreasing in ρ_1 (with the inequality strict when G is strictly increasing in ρ_1). When, instead, G is strictly decreasing in ρ_1 , then $\Gamma_1^{UEC}(\rho, \hat{\rho}) < 0$.

Next, consider ID. Using again the formulas for the functions $V_1(\rho'_1; \rho)$ and $V_2(\rho'_2; \rho)$, we have that

$$\begin{aligned} \Gamma_1^{ID}(\rho, \hat{\rho}) &= \frac{M(\hat{\rho}_1, \hat{\rho}_2)h_\varepsilon}{M^2(\hat{\rho}_1, \hat{\rho}_2)h_\varepsilon + R^2(\hat{\rho}_1, \hat{\rho}_2)h_\omega} [M(\hat{\rho}_1, \hat{\rho}_2)\omega_0 + A(\hat{\rho}_1, \hat{\rho}_2) - M(\rho_1, \hat{\rho}_2)\omega_0 - A(\rho_1, \hat{\rho}_2)] \\ &\quad - \frac{M(\hat{\rho}_1, \rho_2)h_\varepsilon}{M^2(\hat{\rho}_1, \rho_2)h_\varepsilon + R^2(\hat{\rho}_1, \rho_2)h_\omega} [M(\hat{\rho}_1, \rho_2)\omega_0 + A(\hat{\rho}_1, \rho_2) - M(\rho_1, \rho_2)\omega_0 - A(\rho_1, \rho_2)] \end{aligned}$$

and

$$\begin{aligned} \Gamma_2^{ID}(\rho, \hat{\rho}) &= -\frac{M^2(\hat{\rho}_1, \hat{\rho}_2)R^2(\hat{\rho}_1, \hat{\rho}_2)h_\varepsilon + R^4(\hat{\rho}_1, \hat{\rho}_2)h_\omega}{(M^2(\hat{\rho}_1, \hat{\rho}_2)h_\varepsilon + R^2(\hat{\rho}_1, \hat{\rho}_2)h_\omega)^2} + \frac{M^2(\hat{\rho}_1, \rho_2)R^2(\hat{\rho}_1, \rho_2)h_\varepsilon + R^4(\hat{\rho}_1, \rho_2)h_\omega}{(M^2(\hat{\rho}_1, \rho_2)h_\varepsilon + R^2(\hat{\rho}_1, \rho_2)h_\omega)^2} \\ &\quad + \frac{M^2(\rho_1, \hat{\rho}_2)R^2(\rho_1, \hat{\rho}_2)h_\varepsilon + R^4(\rho_1, \hat{\rho}_2)h_\omega}{(M^2(\rho_1, \hat{\rho}_2)h_\varepsilon + R^2(\rho_1, \hat{\rho}_2)h_\omega)^2} - \frac{M^2(\rho_1, \rho_2)R^2(\rho_1, \rho_2)h_\varepsilon + R^4(\rho_1, \rho_2)h_\omega}{(M^2(\rho_1, \rho_2)h_\varepsilon + R^2(\rho_1, \rho_2)h_\omega)^2}. \end{aligned}$$

Hence, for any pair of profiles $\rho = (\rho_1, \rho_2)$ and $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ with $\hat{\rho}_i \geq \rho_i$, $i = 1, 2$, we have that $\Gamma_1^{ID}(\rho, \hat{\rho}) \geq 0$ if the function G is increasing in ρ_2 and the function L is supermodular (strictly, if G is strictly increasing and/or if L is strictly supermodular). Similarly, for the same profiles, we have that $\Gamma_2^{ID}(\rho, \hat{\rho}) \geq 0$ if the function Z is submodular (with the inequality strict, if Z is strictly submodular). Q.E.D.

Proof of Proposition 11. We start by characterizing the ex-ante expected gross payoffs in the more general version of the game where cognition has both a self-directed and a manipulative component. We then specialize the analysis to the case of counter-espionage of Subsection ?? in which cognition is purely manipulative and establish the properties in the proposition.

When player i expects the opponent to engage in cognition $\rho_j = (l_j, t_j)$ and chooses cognition $\rho_i = (l_i, t_i)$ then, given $s_i = (s_i^P, s_i^S)$, player i 's expectation of the fundamental variable ω is equal to

$$\mathbb{E}[\omega | s_i] = \frac{1}{1 + h_i^S} s_i^P + \frac{h_i^S}{1 + h_i^S} s_i^S$$

where

$$h_i^S \equiv [\text{var}(\varepsilon_j + \gamma_j + \phi_i)]^{-1} = \frac{r_i}{1 + r_i}$$

is the precision of the total noise $\varepsilon_j + \gamma_j + \phi_i$ in player i 's secondary signal, with $r_i \equiv [\text{var}(\gamma_j + \phi_i)]^{-1} = \left[\frac{1}{t_j} + \frac{1}{l_i}\right]^{-1} = \frac{t_j l_i}{t_j + l_i}$.

It is then easy to see that the structure of the best responses, as well as that of the equilibrium strategies for the stage-2 game and of the ex-ante expected gross payoffs is identical to the one in Subsection (4.2.2), but with r_i replacing ρ_i in each of the relevant formulas. This observation establishes the various properties in the general case where cognition has both a self-directed and a manipulative dimension (the derivations are omitted for brevity).

Next, let

$$D \equiv \left[V_i(\hat{\rho}_i; \rho') - V_i(\rho_i; \rho') \right] - \left[V_i(\hat{\rho}_i; \rho'') - V_i(\rho_i; \rho'') \right].$$

Using the characterization of the V_i functions in ((18)), after some algebra, we have that

$$\begin{aligned} D &= -2 \left(m_i^{\rho_i; \rho'} - m_i^{\rho_i; \rho''} \right) \left[m_i^{\rho_i; \rho'} + m_i^{\rho_i; \rho''} - 2 + (1 - \beta) \right] \\ &\quad 2 \left(m_i^{\hat{\rho}_i; \rho'} - m_i^{\hat{\rho}_i; \rho''} \right) \left[m_i^{\hat{\rho}_i; \rho'} + m_i^{\hat{\rho}_i; \rho''} - 2 + (1 - \beta) \right] \\ &\quad + 4\beta m_j^{\rho''} \left(m_i^{\hat{\rho}_i; \rho''} - m_i^{\rho_i; \rho''} \right) - 4\beta m_j^{\rho'} \left(m_i^{\hat{\rho}_i; \rho'} - m_i^{\rho_i; \rho'} \right) \\ &\quad + \left(1 - m_i^{\rho_i; \rho'} \right)^2 \frac{1}{r_i^{\rho_i; \rho'_j}} + \beta \left(1 - m_j^{\rho'} \right)^2 \frac{1}{r_j^{\rho_i; \rho'_j}} - \left(1 - m_i^{\hat{\rho}_i; \rho'} \right)^2 \frac{1}{r_i^{\hat{\rho}_i; \rho'_j}} - \beta \left(1 - m_j^{\rho'} \right)^2 \frac{1}{r_j^{\hat{\rho}_i; \rho'_j}} \\ &\quad + \left(1 - m_i^{\hat{\rho}_i; \rho''} \right)^2 \frac{1}{r_i^{\hat{\rho}_i; \rho''_j}} + \beta \left(1 - m_j^{\rho''} \right)^2 \frac{1}{r_j^{\rho_i; \rho''_j}} - \left(1 - m_i^{\rho_i; \rho''} \right)^2 \frac{1}{r_i^{\rho_i; \rho''_j}} - \beta \left(1 - m_j^{\rho''} \right)^2 \frac{1}{r_j^{\rho_i; \rho''_j}}. \end{aligned}$$

Now consider the special case where cognition is purely manipulative so that

$$r_i^{\rho'_i; \rho_j} = t_j \quad \text{and} \quad r_j^{\rho'_i; \rho_j} = t'_i.$$

Let $\rho_i = r_j$, $i, j = 1, 2$, $j \neq i$. As explained in the main text, in this case, because player j cannot observe deviations in cognition by player i and because such deviations do not affect the precision of player i 's own information, for any $\rho = (\rho_1, \rho_2)$,

$$m_i^{\rho'_i; \rho} = \frac{1 + \rho_j(1 + \beta) - 2\beta\rho_j m_j^{\rho}}{1 + 2\rho_j}$$

is invariant in player i 's actual cognition ρ'_i . This property in turn implies that D reduces to

$$D = \beta \left[1 - m_j^{\rho'} - \left(1 - m_j^{\rho''} \right) \right] \left(2 - m_j^{\rho'} - m_j^{\rho''} \right) \left(\frac{1}{\rho_i} - \frac{1}{\hat{\rho}_i} \right).$$

Hence we have that

$$D \stackrel{\text{sgn}}{\equiv} \beta \left[1 - m_j^{\rho'} - \left(1 - m_j^{\rho''} \right) \right] \left(\frac{1}{\rho_i} - \frac{1}{\hat{\rho}_i} \right).$$

Next, use the fact that, for any ρ

$$1 - m_j^{\rho} = \rho_i(1 + \beta) \left\{ \frac{1 + 2\rho_j(1 - \beta)}{1 + 2(\rho_i + \rho_j) + 4\rho_i\rho_j(1 - \beta^2)} \right\}$$

to note that $D \stackrel{\text{sgn}}{\equiv} \beta(\hat{\rho}_i - \rho_i)H$, where

$$\begin{aligned} H &\equiv \rho'_i - \rho''_i + 2\rho'_i\rho'_j(1 - \beta) + 2\rho'_i\rho''_j - 2\rho''_i\rho''_j(1 - \beta) - 2\rho''_i\rho'_j \\ &\quad + 4\beta\rho'_i\rho''_i \left(\rho''_j - \rho'_j \right) (1 - \beta) + 4(\rho'_i - \rho''_i) \rho'_j\rho''_j(1 - \beta). \end{aligned}$$

We are now ready to establish the results in the proposition. To see whether this game satisfies UEC, take $\rho' = (\hat{\rho}_i, \rho_j)$ and $\rho'' = (\rho_i, \rho_j)$. We then have that

$$H = (\hat{\rho}_i - \rho_i) \left(1 + 4\rho_j^2(1 - \beta) + 2\rho_j(1 - \beta) + 2\rho_j \right)$$

Hence, $D \stackrel{\text{sgn}}{\equiv} \beta$.

Next, to see whether this game satisfies ID, take $\rho' = (\hat{\rho}_i, \hat{\rho}_j)$ and $\rho'' = (\hat{\rho}_i, \rho_j)$. We then have that

$$H = -2\beta\hat{\rho}_i(\hat{\rho}_j - \rho_j)(1 + 2\rho_j(1 - \beta))$$

We conclude that $D \stackrel{\text{sgn}}{\equiv} -(\hat{\rho}_i - \rho_i)(\hat{\rho}_j - \rho_j)$. Q.E.D.