

Expectation Conformity in Strategic Cognition*

Alessandro Pavan[†] and Jean Tirole[‡]

December 10, 2022

Abstract

In “cognitive games,” players can influence their understanding of the game (self-directed cognition) and/or their opponents’ understanding (manipulative cognition). The concept of expectation conformity sheds light on cognitive postures and their sensitivity to the type of strategic interaction (e.g., complements vs substitutes). Self-fulfilling cognition never arises in constant-sum games but plays a key role in many non-constant-sum games, both when cognition is self-directed and takes the form of “sparsity,” noisy information acquisition, “espionage,” or level-k reasoning and when it is manipulative and takes the form of framing, signal jamming, noisy disclosures, and counter-intelligence.

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

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

*For useful comments and suggestions, we thank seminar participants at various workshops and institutions where the paper was presented. A special thank is to Robert Gibbons for very detailed feedback. Financial support from the European Union’s Horizon 2020 research and innovation program (Grant Agreement no. 669217) is gratefully acknowledged. Matteo Camboni provided 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 the 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 à 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-espionage), 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 some of the models considered in the career-concerns literature (see, e.g., 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 traps obtain 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". 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 The reason for considering this application is twofold. First, sparsity is believed to play an important role in many socio-economic environments (see Gabaix, 2014, for an introduction to the concept of sparsity and for a discussion of the role that it plays in many decision-theoretic settings). Second, it permits us to illustrate the importance that expectation conformity plays in settings in which limited cognition leads to behavioral patterns that may appear inconsistent with full rationality. Sparsity is often intended to capture a certain form of bounded rationality. We show that its key features are also consistent with a certain representation where all players are fully rational.

In strategic settings, 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

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

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 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 players learn about other players' beliefs. The reason for considering this application is that it permits us to capture situations in which cognition takes the form of putting oneself in another player's shoes. Examples of such situations include instances in which players invest to interpret other players' culture, cognitive traits, and other dimensions of their personality that are responsible for their view of the game, but also instances in which players spy on other players' information to better predict their ultimate behavior, as in certain industrial-espionage games (see the recent article "Corporate espionage is entering a new era," Economist magazine, June 2022).

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.

Finally, we consider games in which players invest in making themselves understood or in preventing others from understanding them. Examples include situations in which players share their information with others (as in certain noisy communication games, but also in games in which players disclose hard information to others), or take actions meant to obstruct others from spying on their own information (as when firms invest in counter-espionage to make it difficult for rival firms to scoop on their research output). The reason for considering such situations is that they permit us to illustrate the role that expectation conformity plays in settings in which cognition is *manipulative* (that is, it affects other players' understanding of the game).

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-espionage to prevent others from spying on them. This is true irrespective of whether downstream actions are complements or substitutes.

In an Online Supplement, we extend the analysis to a broader class of cognitive situations. First, we consider games in which players receive noisy additive signals about an exogenous payoff state, where cognition is self-directed, with each player's cognition determining the precision of her own signal (the case most studied in the literature on information acquisition in games). Next, we consider a different class of games of manipulative cognition in which some players choose frames to influence other players' recollection of information previously received. We also consider a class of

games of manipulative cognition that generalizes Holmström (1999) signal-jamming model of career concerns. In all of these settings, cognition takes the form of the choice of an information structure (for oneself, when self-directed, or for other players, when manipulative). The last section of the Online Supplement considers instead situations in which payoffs are commonly known and cognition 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.

Summarizing, we view the paper’s contribution as twofold. First, in showing how the choice of the cognitive posture reflects the players’ beliefs over (a) the cognition of others and (b) other players’ expectations of their own cognition; EC, and its decomposition into UEC and ID, helps connecting the determinacy of equilibria, as well as the possibility of cognitive traps, to the type of strategic interactions in the primitive game, but also facilitates comparative statics. Second, the paper provides a unifying perspective over apparently different forces that shape the cognitive choices. By bringing distinct phenomena under the same conceptual umbrella, the analysis helps identifying common themes but also isolating 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 contains the description of cognitive games. It also contains the definition of EC, UEC and ID, and the results relating these concepts to the determinacy of equilibria and to constant-sum games. Section 3 contains the analysis of games in which cognition is self-directed and takes the form of sparsity. Section 4 contains the results for games in which cognition takes the form of noisy information acquisition about other players’ beliefs (espionage) as well as the case in which cognition is manipulative and takes the form of noisy information sharing and counter-espionage. Section 5 discusses the role that EC, and its decomposition into UEC and ID, plays in various other settings formally analyzed in the Online Supplement. Section 6 concludes. 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, 2021) 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).² 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, 2022a,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.³ Instead,

²See also Bergemann and Morris (2013) for an analysis of robust predictions in these games.

³Denti (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

Dewatripont and Tirole (2005), Che and Kartik (2009), Calvo-Armengol et al. (2015), Sethi and Yildiz (2016, 2018), Kozlovskaya (2018), Carroll (2019), 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.⁴

A second strand is the literature on sparsity and endogenous depth of reasoning. See Gabaix (2014) for the former and Alaoui and Penta (2016, 2022) and Alaoui et al. (2020) for the latter.⁵ 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 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 Hörner 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, Bénabou, 2013, Gottlieb 2014a,b, and Salant and Siegel, 2018).

The strands of the literature mentioned above are too vast to be described in more detail 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.

observe other agents' actions and prices with endogenous noise.

⁴Compared 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. (2022), Hebert and La'O (2020), and Hebert and Woodford (2020) for related models of rational inattention with alternative cost functionals.

⁵At a broad level, the paper is also related to the literature studying how players conceptualize strategic situations. See, for example, Heller and Winter (2016), and Gibbons et al. (2021).

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 .⁶ 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}) , where $a_{-i} \equiv (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n) \in \prod_{j \neq i} A_j$.

Cognition. The players have a common prior F on the state space Ω .⁷ Prior to playing the stage-2 game, the players engage in cognition at stage 1, resulting in a vector of information structures. We index player i ’s cognition 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 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 $Q(s|\omega, \rho)$.¹⁰

⁶Such payoffs can, but need not, be interpreted as expected utility.

⁷Che and Kartik (2009) by contrast look at incentives to acquire information in an environment with heterogeneous priors. Our definition of expectation conformity and most of our results do not hinge on the players sharing a common prior. However, the exposition is facilitated by such an assumption.

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 5.3 and formally treated in the Online Appendix.

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

¹⁰Note that the realization s of the jointly controlled experiment Q may carry information not only about the state, ω , but also about the cognitive postures, ρ . Also note that, at this level of abstraction, it is irrelevant whether each ρ_i is interpreted as a deterministic or a stochastic choice (i.e., a mixture over a set of cognitive postures). Likewise, the analysis does not distinguish between the case in which cognition is static and the one in which it is dynamic. The latter case is captured by interpreting each ρ_i as a complete dynamic cognitive plan specifying how the player responds dynamically to the (unmodelled) gradual resolution of uncertainty leading to the determination of the signal s_i (with the latter interpreted as the collection of intermediate pieces of information received by the player over time, prior to

We say that cognition is *self-directed* when the distribution from which each player's information is drawn is not affected by 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) , $Q(s|\omega, \rho) = \bigoplus_{i \in I} Q_i(s_i|\omega, \rho_i)$. Note that, when cognition is self-directed, fixing the expectations of player i 's opponents over player i 's cognition, the actual choice of cognition by player i cannot be detected by player i 's opponents from the observation of their signals $s_{-i} \equiv (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$. The opposite is not true. Cognition can be manipulative without any player being able to detect deviations by the other players; this happens when player i 's cognition affects the distribution from which the other players' signals s_{-i} is drawn but not its support.

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 's signal under ρ'_i is Blackwell more informative than under ρ_i .¹¹ 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.¹²

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 and reduced-form payoffs. 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 $s_i \in S_i$, a mixed action $\sigma_i(s_i) \in \Delta(A_i)$.¹³

When players engage in cognition ρ and play in the downstream game according to σ , player i expects a gross payoff 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 ρ of focal cognitive choices, let σ^ρ denote the stage-2 strategy profile that the players expect will be used in the downstream game when cognition in the upstream game is commonly expected to be equal to ρ . In most applications of interest, σ^ρ can be taken to be the stage-2

the determination of the downstream action α_i).

¹¹When signals take the form of partitions of Ω , $\rho_i < \rho'_i$ means that ρ'_i is finer than ρ_i . The reason why we model information structures as a jointly controlled experiment instead of a collection of partitions is that the latter modellization is not fully flexible when cognition is manipulative.

¹²This 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 when information acquisition is covert: the unbundled information structure is simply irrelevant and can be assumed to be infinitely costly (this is not the case with overt information acquisition — a player may suffer when other players know that he has more information).

¹³Note that, when useful to highlight it, the dependence of player i 's behavior in the downstream game on his own upstream cognition ρ_i can always be captured by letting s_i contain ρ_i .

continuation equilibrium associated with the cognitive profile ρ .¹⁴ More broadly, σ^ρ should be interpreted as the strategy profile that the players expect will be used in the downstream game when the players engage in cognition ρ .

Next, for any $\rho = (\rho_i, \rho_{-i})$ and ρ'_i , let

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

denote the maximal gross payoff that player i can obtain by selecting cognition ρ'_i and then adjusting his stage-2 behavior optimally, when all players other than i choose cognition ρ_{-i} and play according to the strategies σ_{-i}^ρ in the downstream game. In particular, when $\rho'_i = \rho_i$,

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

denotes the gross payoff that player i obtains by selecting the cognition ρ_i specified in the profile $\rho = (\rho_i, \rho_{-i})$ and then playing according to σ_i^ρ (with all the other players selecting cognition ρ_{-i} and playing according to σ_{-i}^ρ). In other words, for any $(\rho'_i; (\rho_i, \rho_{-i}))$, the term ρ'_i at the left of the semicolon will always correspond to the actual choice by player i , whereas the component ρ_i in the profile $\rho = (\rho_i, \rho_{-i})$ to the right of the semicolon will always correspond to the cognitive level expected from player i .

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 $\sigma_{-i}^{\rho_i, \rho_{-i}}$.

Our notion of expectation conformity (formally introduced in Section 2.2 below), as well as the use we make of it to shed light on a player's incentives to choose one cognitive posture over another, does not require to confine attention to equilibrium analysis. However, in many applications, one may be interested in using the results to identify the cognitive profiles that are sustained in equilibrium. In this case, we will say that the cognitive profile ρ is part of a Perfect Bayesian Equilibrium (in short, is an “equilibrium” cognitive profile) if, for any player i and any ρ'_i ,

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

with σ^ρ describing the corresponding downstream strategy profile. To ease the exposition, we will omit the specification of the beliefs associated with each equilibrium cognitive profile. Although not necessary, to fix ideas, we suggest the reader interprets the game described above as one in which cognition is covert (i.e., the distribution from which player i 's signal s_i is drawn may depend on his opponents' cognitive choices – in case cognition is manipulative – but the support of such

¹⁴In case there are multiple continuation equilibria, we assume that some selection is in place. Also note that, when the realization of the signals s carries information about ρ , σ^ρ , being a collection of complete plans of action, also specifies downstream actions for signal realizations that need not be in the support of $Q(\cdot|\omega, \rho)$. Depending on the solution concept of choice, σ^ρ may then be required to satisfy appropriate perfection requirements.

a distribution is invariant in the choices by player i 's opponents). In this case, beliefs are always pinned down by Bayes rule and, for any ρ, σ^ρ is a Bayes-Nash equilibrium in the hypothetical game in which players' cognition is fixed at ρ .

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 player i with respect to cognitive profiles ρ and $\hat{\rho}$ if

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

Hereafter, when we say that expectation conformity holds for profiles ρ and $\hat{\rho}$ we mean for all players.

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' cognitive choices at 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 captures the impact of a variation in other players' cognition on player i 's cognitive choice, holding fixed the other players' expectations of player i 's cognition.

When expectation conformity holds for player i with respect to 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 increasing differences, 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

We now show how expectation conformity relates to the determinacy of equilibria. For this purpose, we now interpret σ^ρ as the continuation equilibrium associated with the cognitive profile ρ . We also assume that, for any ρ , a continuation equilibrium σ^ρ exists. 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 player i when acquired covertly (i.e., when variations in ρ_i cannot be detected by player i 's opponents), 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.

¹⁵Clearly, this property need not hold when players are boundedly rational; see the discussion of the level-k model in the Online Supplement.

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

2.4 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 2 (constant-sum games). *Two-person constant-sum games satisfy the following property, for all $(\rho, \hat{\rho})$:*

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

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 $(\rho, \hat{\rho})$,

$$\sum_i \left\{ \left[U_i(\sigma_i^{\hat{\rho}}, \sigma_j^{\hat{\rho}}; \hat{\rho}_i, \hat{\rho}_j) - U_i(\sigma_i^{\rho}, \sigma_j^{\rho}; \rho_i, \rho_j) \right] - \left[U_i(\sigma_i^{\hat{\rho}}, \sigma_j^{\rho}; \hat{\rho}_i, \rho_j) - U_i(\sigma_i^{\rho}, \sigma_j^{\rho}; \rho_i, \rho_j) \right] \right\} = 0. \quad (2)$$

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

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

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 (2) implies that

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

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

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

$$V_i(\hat{\rho}_i; \hat{\rho}) - V_i(\rho_i; \hat{\rho}) = C_i(\hat{\rho}_i) - C_i(\rho_i) = V_i(\hat{\rho}_i; \rho) - V_i(\rho_i; \rho). \quad (3)$$

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.¹⁷

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.

3 Sparsity in games

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) and takes the form of sparsity. As anticipated in the Introduction, the latter is typically interpreted as a form of bounded rationality (see, e.g., Gabaix, 2014). The formalism below allows us to capture many of the key features of sparsity while retaining the convenience of a framework with fully rational players.

The players’ payoffs are given by

$$u_i(a_i, a_{-i}, \omega) = -(1 - \beta)(a_i - g(\omega))^2 - \beta(a_i - \bar{a}_{-i})^2 + \psi(a_{-i}, \omega), \quad (4)$$

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 $K \in \mathbb{N} \cup \{+\infty\}$, $\bar{a}_{-i} \equiv \sum_{j \neq i} a_j / (n - 1)$, $g : \mathbb{R}^K \rightarrow \mathbb{R}$, 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} 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*.¹⁸ Finally, the function

¹⁷To illustrate this point in the simplest possible terms, suppose i ’s payoff is $(a_i - a_j)\omega$, for $i, j = 1, 2$, $j \neq i$. Further assume that $A_i = \{1, -1\}$, for $i = 1, 2$, 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 . When, instead, $\rho_i = \emptyset$, player i receives no information about ω . Let $C_i(1) = 1$ and $C_i(\emptyset) = 0$. The game admits multiple pure-strategy equilibria. In one, both players learn the state and then match the state with their action. In another one, neither player learns the state and then selects action $a_i = 1$ with probability one, $i = 1, 2$. The equilibrium payoffs under the first equilibrium are equal to -1 for both players, whereas the equilibrium payoffs under the second equilibrium are equal to 0.

¹⁸The formulas below extend to the case in which $\beta = 0$, but this case is not interesting, for it corresponds to a non-strategic situation.

$\psi(a_{-i}, \omega)$ summarizes various external effects that matter for payoffs but do not play any role for the selection of the individual best responses.

The statistics $g(\omega)$ takes the form $g(\omega) = (1 + \sum_{k=1}^K \omega^k)/(1 - \beta)$. 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 . As shown below, the combination of these assumptions with the quadratic payoffs in (4) is what permits us to capture the essence of sparsity in an otherwise fully-rational model.

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 ($\omega^1, \dots, \omega^{k-1}$). This assumption is not really needed for our analysis of the role that expectation conformity plays in such games. It has the advantage though of permitting us to interpret each player's cognition with her "depth of knowledge."¹⁹ That is, we let each player's cognition coincide with the number of dimensions $\rho_i \in \mathbb{N}$ of the state ω the player learns about.²⁰

A player who decides to learn ρ_i dimensions of the state then reasons "as if" the remaining $K - \rho_i$ dimensions did not exist (that is, as if the state had only ρ_i dimensions). This is one of the key distinctive features of sparsity. A second key feature (which the above formalization captures) 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 did, 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 $s_i \equiv (\omega^1, \dots, \omega^{\rho_i})$ be the subset of the state ω explored by player i . Using the general notation of Section 2, we then have that, for any (ω, ρ) , $Q(s|\omega, \rho)$ is a Dirac measure assigning probability one to the signal vector $s = (s_1, \dots, s_n)$, where n is the number of players.

For simplicity, assume that $n = 2$ and, without loss of generality, let player 1 be the player with the lowest depth of knowledge, that is, for whom $\rho_1 \leq \rho_2$.²¹ Given the cognitive profile $\rho = (\rho_1, \rho_2)$, in the stage-2 game, there exists a unique continuation equilibrium and is such that, for any ω , the

¹⁹In the absence of such an assumption (that is, when cognition is not ordered, so that a player can learn dimension k without having learned all previous dimensions), cognition ρ_i takes the form of a vector of zeros and ones, where each entry indexes whether that dimension has been explored. The analysis below can also easily accommodate for the possibility of "clusters," with dimensions correlated within clusters but independent across clusters (say, with each dimension affected by a cluster-specific shock plus a dimension-specific shock).

²⁰As discussed below, this assumption can also be micro-founded. For example, when the costs C_i are separable, one can order the dimensions so that the benefit of exploring each dimension relative to its cost is decreasing in k . When $\beta \geq 0$, such an assumption suffices to guarantee that the equilibria of the game in which the players must explore the dimensions of the state in an increasing sequence are also equilibria in a game in which they can choose the order of exploration of their choice. When, instead, $\beta < 0$, the same conclusion holds, provided that, in addition, one assumes that the benefit of exploring each dimension, relative to its cost, declines sufficiently fast with k for the players to find it suboptimal to explore dimensions not sequentially as a way of differentiating their action from their opponents.

²¹The 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 for the $n > 2$ case coincide with those for the $n = 2$ case.

two players' equilibrium strategies are Dirac distributions assigning probability one to the actions²²

$$a_1^\rho(s_1) = \frac{1 + \sum_{k=1}^{\rho_1} \omega^k}{1 - \beta}$$

and²³

$$a_2^\rho(s_2) = \frac{1 + \sum_{k=1}^{\rho_1} \omega^k}{1 - \beta} + \sum_{k=\rho_1+1}^{\rho_2} \omega^k.$$

The equilibrium actions thus reflect the property anticipated above that each player acts “as if” each neglected dimension $k > \rho_i$ did not exist. In particular, the equilibrium action of the player who is 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 $a_1^{\rho_1, \rho_1}(s_1)$ that the player would choose if the state was commonly known to have only ρ_1 dimensions, with the dimensions $(\omega^1, \dots, \omega^{\rho_1})$ commonly known to both players. Similarly, the stage-2 equilibrium action for the player who is ahead in the exploration of the state (player 2) coincides with the equilibrium action of the player who is behind, $a_1^\rho(s_1)$, 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. The action $a_2^\rho(s_2)$ also coincides with the player's equilibrium action when she expects the state to have only ρ_2 dimensions and her opponent to explore only the first $\rho_1 < \rho_2$ dimensions.

When both players expect cognition $\rho = (\rho_1, \rho_2)$, their ex-ante expected payoffs (gross of the cognitive costs, C_i , and of the expectation of the terms $(1 - \beta)g(\omega)^2 + \beta \bar{a}_{-i}^2 + \psi(a_{-i}, \omega)$ that do not interact with the players' best responses) are then equal to (see the proof of Proposition 3 in the Appendix for details)

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

and

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

Note that an implication of the assumed structure is that the various shocks ω^k impact the players' best responses and payoffs separately. To gather more intuition for the above results (but also for some of the properties derived below) consider for a moment a simplified version of the above problem in which $K = 1$ (that is, the state is unidimensional). Abusing notation, then drop the superscript “1” from the state ω and let σ^2 denote the variance of ω . Furthermore, let $\rho = (\rho_i, \rho_j)$ denote an arbitrary cognitive profile that is commonly expected by both players, and abandon for a moment the convention that player 1 is behind in the exploration of the state, so as to treat the two players symmetrically.

In this simplified problem, $\rho_i \in \{0, 1\}$, with $\rho_i = 1$ denoting the decision to learn the state, and $\rho_i = 0$ the decision to remain uninformed. When $\rho_i = 1$, player i receives the signal $s_i = \omega$ with certainty, whereas, when $\rho_i = 0$, she receives the null signal $s_i = \emptyset$ with certainty.

²²The 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 3.

²³One 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.

Equilibrium cognition: $\rho = (\rho_i, \rho_j)$	Shared knowledge: (1, 1)	Shared ignorance: (0, 0)	Asymmetric knowledge	
			i informed: (1, 0)	i uninformed: (0, 1)
Impulse response, I_i^ρ	$\frac{1}{1-\beta}$	1	1	$1 + \beta$
Value of information, \mathbb{V}_i^ρ	$\frac{\sigma^2}{(1-\beta)^2}$	σ^2	σ^2	$(1 + \beta)^2 \sigma^2$

Table 1: Value of information as a function of ρ

In the stage-2 game, player i 's optimal action is then given by

$$a_i = 1 + \mathbb{E}[\omega | \rho_i, s_i] + \beta \mathbb{E}[a_j | \rho_i, s_i],$$

where $\mathbb{E}[\omega | \rho_i, s_i] = s_i = \omega$ if $\rho_i = 1$ (that is, if player i learned the state) and $\mathbb{E}[\omega | \rho_i, s_i] = 0$ if $\rho_i = 0$ (in which case $s_i = \emptyset$).

Paralleling the analysis also for player j , we then have that, when the two players are expected to engage in cognition $\rho = (\rho_i, \rho_j)$ and, instead, player i chooses cognition ρ'_i , in the stage-2 game, player i then optimally chooses an action equal to

$$a_i = \frac{1}{1-\beta} + \rho'_i I_i^\rho \omega$$

where

$$I_i^\rho \equiv \frac{1 + \beta \rho_j}{(1 - \beta^2 \rho_i \rho_j)}$$

denotes the *impulse response* of player i 's action to the state, ω , when player i chooses to learn the state (i.e., $\rho'_i = 1$) and the two players are expected to engage in cognition $\rho = (\rho_i, \rho_j)$. Such an impulse response naturally reflects the following properties: (a) when player j is not expected to learn the state (i.e., $\rho_j = 0$), the value to player i from learning the state comes entirely from being able to respond to it, with no effect on her ability to respond to the other player's action (the impulse response is then equal to $I_i^\rho = 1$, irrespectively of whether $\rho_i = 0$, or $\rho_i = 1$); (b) when, instead, player j is the only player expected to learn the state (i.e., $(\rho_i, \rho_j) = (0, 1)$), the value to player i from learning the state also accounts for the improvement in her ability to respond to variations in her opponent's action but without the opponent accommodating for the unanticipated joint learning (the impulse response is then equal to $I_i^\rho = 1 + \beta$ which is higher than 1 when action are complements, and lower than 1 when they are substitutes); (c) finally, when both players are expected to learn the state (i.e., $\rho = (1, 1)$), the impulse response is equal to $I_i^\rho = 1/(1 - \beta)$, which is higher than $1 + \beta$ no matter whether actions are complements or substitutes, reflecting the benefit that player i derives from player j responding to the state accounting for joint learning.

We can then use the above properties to compute player i 's value of information

$$\mathbb{V}_i^\rho \equiv V_i(1; \rho) - V_i(0; \rho)$$

as a function of the cognitive profile ρ expected from the two players. The results are summarized in Table 1, which represents the impact of agent i 's selection of cognition $\rho'_i = 1$ instead of $\rho'_i = 0$ on the player's behavior and payoff.

Hence, under strategic complements ($\beta > 0$), $\mathbb{V}_i^{(1,1)} > \mathbb{V}_i^{(0,1)} > \mathbb{V}_i^{(1,0)} = \mathbb{V}_i^{(0,0)}$. The value of learning the state is the highest when both players are expected to learn the state ($\rho = (1,1)$). Learning the state jointly with the opponent but without the opponent expecting i to learn the state ($\rho = (0,1)$) is less valuable but still superior to being the sole learner of the state ($\rho = (1,0)$ or $\rho = (0,0)$). Under strategic substitutes, instead, $\mathbb{V}_i^{(1,0)} = \mathbb{V}_i^{(0,0)} > \mathbb{V}_i^{(1,1)} > \mathbb{V}_i^{(0,1)}$. The value of learning the state is highest when the other player does not learn the state ($\rho = (1,0)$ or $\rho = (0,0)$). Conditional on the other player learning the state, player i is better off if player j anticipates joint learning ($\rho = (1,1)$) than if she doesn't ($\rho = (0,1)$), for j “makes room” to i by reducing her response to the state, when expecting joint learning.

Returning to the case in which $K > 1$, we can then use the above properties to establish the role of expectation conformity in these games.

Proposition 3 (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 $\rho = (\rho_1, \rho_2)$ for which player 2 is always ahead in the exploration of the state (i.e., $\hat{\rho}_2 > \rho_2 \geq \hat{\rho}_1 > \rho_1$), possibly reflecting major differences in the cost of cognition. As anticipated above, for the player who is ahead in the exploration of the state (player 2), the value of exploring more dimensions is the same no matter whether the opponent (player 1) expects player 2 to explore $\hat{\rho}_2 \geq \rho_1$ or $\rho_2 \geq \rho_1$ dimensions. This is because player 1’s stage-2 action is the same under both expectations (recall that a key feature of the sparsity model is that the behavior of the player who is behind in the exploration of the state is invariant in how far ahead she expects the opponent to be in the exploration of the state). For player 1, instead, the value of exploring more dimensions is larger when the opponent expects her to explore more dimensions. As explained above, this is because player 2 changes her response to the extra dimensions ($\rho_1 + 1, \dots, \hat{\rho}_1$) explored by player 1, when she expects player 1 to learn such 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).

That, for such profiles, ID holds as an equality for both players reflects the separability of the expected payoffs in the explored dimensions. For player 1, the value of expanding her cognition from ρ_1 to $\hat{\rho}_1 \leq \rho_2 < \hat{\rho}_2$ is invariant in how far ahead player 2 is in the exploration of the state, reflecting the fact that player 1 treats all dimensions that she does not explore as if they did not exist. Likewise, the value that player 2 assigns to expanding her cognition from ρ_2 to $\hat{\rho}_2$ is invariant to whether she expects player 1 cognition to be $\hat{\rho}_1$ or cognition ρ_1 for, in either case, player 2 does not expect player 1 to respond to the dimensions ($\rho_2 + 1, \dots, \hat{\rho}_2$).

Next, consider profiles $\hat{\rho} = (\hat{\rho}_1, \hat{\rho}_2)$ and $\rho = (\rho_1, \rho_2)$ such that $\hat{\rho}_2 = \hat{\rho}_1 > \rho_2 = \rho_1$. The choice of such profiles is motivated by the interest in understanding whether multiple symmetric (pure-

strategy) equilibria are possible in these games. When a player expands her cognition starting from the cognitive level of her opponent, the value she assigns to the expansion is the same no matter whether she expects the opponent to anticipate the expansion. This is because the opponent does not respond to any dimension that she does not explore and responds to the explored dimensions as if the unexplored ones did not exist. Hence, UEC holds as an equality for these profiles. Fixing the opponent's expectations about her own cognition, we have that the value that player i assigns to exploring more dimensions is larger when the opponent also explores them and downstream actions are strategic complements, whereas it is lower when the opponent also explores the extra dimensions and downstream actions are strategic substitutes. This is because, as explained above, the value of joint learning is higher than the value of sole learning when actions are complements, whereas the opposite is true when actions are substitutes.

When combined with Proposition 1, the results in part (a) of Proposition 3 thus suggests that, when the players face asymmetric cost functions, these games may admit multiple asymmetric (pure-strategy) equilibria, both in the complements and in the substitutes case. The results in part (b), instead, suggest that, when the cost functions are symmetric, these games are likely to feature a unique symmetric (pure-strategy) equilibrium when actions are substitutes, and multiple symmetric (pure-strategy) equilibria when actions are complements. The next two propositions verify these conjectures by considering cost functions that are separable over the explored dimensions. The propositions also provide a more detailed account of the type of multiplicity these games are prone to.

Proposition 4 (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 form $C_i(\rho_i) = \sum_{k=1}^{\rho_i} c_k$, 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 $k^ \in [\underline{k}, \bar{k}(\beta)] \cap \mathbb{N}$ 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 the state ω and on other agents' actions a_{-i} only through the effects that the latter have on individual best responses.²⁴ Then the (pure-strategy) equilibria are Pareto ranked, with the players' net payoff increasing in the equilibrium depth of knowledge k^* .
- All equilibria of the game in which cognition is ordered are also equilibria of the game in which cognition is unordered (i.e., players can explore a dimension k without having explored all

²⁴When payoffs are as in (4), this occurs when $\psi(a_{-i}, \omega) = (1-\beta)g(\omega)^2 + \beta \bar{a}_{-i}^2$. In this case, the term $-(1-\beta)g(\omega)^2$ from the first addendum in (4) is perfectly offset by the corresponding term in ψ and likewise for the externality term $\beta \bar{a}_{-i}^2$ originating in the second addendum in (4). As a result, individual payoffs u_i depend on ω and on other agents' actions a_{-i} only through the effect that the latter terms have on individual best responses.

dimensions $k' < k$). Furthermore, when, for any $k \in [\underline{k}, \bar{k}(\beta)] \cap \mathbb{N}$,

$$\frac{\sigma_{k-1}^2}{c_{k-1}} > \frac{\sigma_k^2}{c_k(1-\beta)^2}, \quad (5)$$

the converse is also true: any equilibrium of the game in which cognition is unordered is also an equilibrium of the game in which cognition is ordered.

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 when both players learn that dimension than when a player is 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 $k^* + 1$ even if such dimension is explored solely (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 the net benefit of learning any dimension $k > \bar{k}(\beta)$ is negative even when learnt jointly with the other player ($\sigma_k^2/(1-\beta)^2 - c_k < 0$ for $k > \bar{k}(\beta)$). That, for any $k \in [\underline{k}, \bar{k}(\beta)] \cap \mathbb{N}$, a symmetric equilibrium with depth of cognition k^* exists in turn follows from the characterization of the payoff functions $V_i(\rho'_i; \rho)$ provided above, along with the fact that σ_k^2/c_k is decreasing in k .

Finally, that all equilibria in the game in which cognition is ordered are also equilibria in the game in which cognition is unordered follows from the fact that (a) the benefit σ_k^2/c_k of learning dimension k , relative to its cost, is decreasing in k , along with (b) the separability of the players' payoffs in the explored dimensions. That the converse is also true when Condition (5) holds for any $k \in [\underline{k}, \bar{k}(\beta)] \cap \mathbb{N}$ follows from the fact that any equilibrium of the game in which cognition is unordered is necessarily symmetric, with the smallest dimension explored greater than \underline{k} and the largest one smaller than $\bar{k}(\beta)$, along with the fact that, by virtue of Condition (5), when exploring dimension $k \in [\underline{k}, \bar{k}(\beta)] \cap \mathbb{N}$ jointly with the other player is profitable, the net benefit of exploring any lower dimension $k' < k$ is strictly positive even if such a dimension is explored solely.

The above results thus imply that, when downstream actions are complements and payoffs are symmetric across players and separable over the dimensions explored, the depth of knowledge is not uniquely pinned down in equilibrium. However, in the absence of external payoff effects, such games never exhibit cognitive traps: the players are better off the larger the depth of knowledge. This is because more knowledge always permits the players to better align their actions with the underlying state.

We now turn to games in which downstream actions are strategic substitutes. Consistently with the convention above, we refer to player 2 as the player going deeper in the exploration of the state.

Proposition 5 (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 $C_i(\rho_i) = \sum_{k=1}^{\rho_i} c_k$, 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*

$$\frac{\sigma_{k^*+1}^2}{c_{k^*+1}} \leq 1 \leq \frac{\sigma_{k^*}^2}{(1-\beta)^2 c_{k^*}}. \quad (6)$$

When a symmetric (pure-strategy) equilibrium exists, it is unique and its cognitive level k^ satisfies (6).*

- *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

$$\frac{\sigma_{\rho_1+1}^2 (1+\beta)^2}{c_{\rho_1+1}} \leq 1 \leq \frac{\sigma_{\rho_1}^2}{(1-\beta)^2 c_{\rho_1}} \quad (7)$$

and

$$\frac{\sigma_{\rho_2+1}^2}{c_{\rho_2+1}} \leq 1 \leq \frac{\sigma_{\rho_2}^2}{c_{\rho_2}}. \quad (8)$$

Player 1's equilibrium payoff is increasing in her 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 (8), is invariant in ρ_2 .²⁵ The sum of the two players' equilibrium payoffs is maximal under the (pure-strategy) equilibrium featuring the lowest cognition for player 1 (the one behind in the exploration of the state).

- *All the equilibria of the game in which cognition is ordered are also equilibria of the game in which cognition is unordered. Furthermore, when, for any $k \in [\underline{k}(\beta), \bar{k}]$,*

$$\frac{\sigma_{k-1}^2 (1+\beta)^2}{c_{k-1}} > \frac{\sigma_k^2}{c_k}, \quad (9)$$

the converse is also true: any equilibrium of the game in which cognition is unordered is also an equilibrium of the game in which cognition is ordered.

That, with substitutes, such games admit at most one symmetric equilibrium follows from the fact that the gross benefit of learning any dimension is largest when learnt solely. Furthermore, when

²⁵This is because multiple solutions to the two inequalities in (8) 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.

a dimension is jointly learnt, the gross benefit of learning it is larger when the opponent expects joint learning than when she expects to be the sole learner.²⁶ The double inequality in Condition (6) then guarantees that, starting from a situation where both players learn the same dimensions, no player has a profitable deviation. Clearly, there is at most one dimension k^* satisfying the double inequality in (6). Hence, there exists 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 (6).

Next, consider asymmetric (pure-strategy) equilibria. In any such equilibrium, the value to player 2 (the one ahead in the exploration of the state) 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 (8). 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 3, observe that the term $\sigma_{\rho_1+1}^2(1+\beta)^2 - c_{\rho_1+1}$ in the left-hand side of (7) is the net benefit of exploring the (ρ_1+1) -th dimension when player 2 also explores such a 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 (7) is the net benefit of exploring the ρ_1 -th dimension when player 2 also explores such a 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.

The ranking of the asymmetric equilibria in terms of the players' payoffs naturally reflects the fact that each player's *equilibrium* payoff is increasing in the number of dimensions explored. Hence, these games too are never conducive to cognitive traps.

That player 1's equilibrium payoff is invariant in player 2's equilibrium depth of knowledge follows from the fact that both players' response to the dimensions learnt jointly is invariant in the number of dimensions explored solely by player 2, along with the separability of the payoffs. That, instead, player 2's equilibrium payoff is decreasing in player 1's depth of knowledge follows from the fact that, with substitutes, learning a dimension solely is always more valuable than learning it jointly. That the sum of the two players' equilibrium payoffs

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

is maximal under the equilibrium featuring the lowest cognition 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 each dimension that player 2 also learns is smaller than the loss $\sigma_k^2 - \sigma_k^2/(1-\beta)^2$ that player 2 incurs when player 1 learns the additional dimension.

Finally, consider the comparison of the equilibria in the game in which cognition is ordered with those in the game in which cognition is unordered. That any equilibrium of the game in which cognition is ordered is also an equilibrium in the game in which cognition is unordered follows from the monotonicity of the relative benefits, σ_k^2/c_k , along with the separability of the payoffs in the dimensions explored. That, under Condition (9), any equilibrium in the game in which cognition is unordered coincides with one of the equilibria in the game in which cognition is ordered follows from the following properties. In any equilibrium of the game in which cognition is unordered, the

²⁶Use Table 1 to verify such properties.

smallest dimension explored is necessarily greater than $\underline{k}(\beta)$ whereas the largest one is necessarily smaller than \bar{k} , exactly as in the game in which cognition is ordered. Condition (9) in turn implies that, when the net value to player i of exploring dimension k is positive, so is the value of exploring any lower dimension $k' < k$. This is so irrespective of whether dimensions k and k' are explored solely by player i or jointly with player j , and irrespective of player j 's expectations about player i 's cognition. Hence, whenever a player explores dimension k , she necessarily explores also any lower dimension.²⁷

4 Learning about others' beliefs : Espionage and counter-espionage

We now move to a class of games, inspired by the literature on noisy communication (see, e.g., Dewatripont and Tirole, 2005, and Calvo-Armengol et al. 2015), in which the players' cognition affects their understanding of other players' beliefs. The stage-2 game has the same structure as in the previous section, with payoffs as in (4), but with $K = 1$ (uni-dimensional state), $g(\omega) = \omega$, and $\psi(\omega, a_{-i}) = 0$. To ease the exposition, assume that the variable ω is drawn from an improper uniform prior over the entire real line. Each player is endowed with an exogenous "primary" signal equal to

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

with ε_i drawn from a standard Normal distribution. Such a signal captures the player's exogenous private information (equivalently, his primitive beliefs). In addition to s_i^P , player i receives a "secondary" signal

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

about the opponent's primary signal, with γ_j drawn from a Normal distribution with mean zero and precision t_j , and with ϕ_i drawn from a Normal distribution with mean zero and precision l_i . The variables $(\omega, \varepsilon_1, \varepsilon_2, \gamma_1, \gamma_2, \phi_1, \phi_2)$ are jointly independent.

A player's cognition affects the precision of her secondary signal (espionage) and/or the precision of her opponent's secondary signal (counter-espionage). Specifically, player i 's cognition is given by $\rho_i = (l_i, t_i)$. The term l_i captures player i 's investment to learn player j 's beliefs (through the precision of the noise ϕ_i in s_i^S). The term t_i , instead, captures player i 's effort to influence player j 's ability to learn player i 's primary information (through the precision of the noise term γ_i in s_j^S). The l_i dimension thus corresponds to the *self-directed* component of player i 's cognition, whereas the t_i dimension corresponds to the *manipulative* component of i 's cognition. A higher t_i represents a larger investment by player i in activities facilitating agent j 's ability to interpret agent i 's primitive information. Alternatively, l_i can be thought of as player i 's investment in espionage, whereas $1/t_i$ as her investment in counter-espionage, that is, in activities that make it more difficult for player j to "spy" on i 's primitive information.

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 either

²⁷Note that Condition (9) also implies that there can be at most one dimension explored solely by one of the two players.

purely self-directed (espionage) or manipulative (counter-espionage), so as to isolate the specific effects of the two cognitive modes.

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. 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 (making oneself understood is costly). If, instead, t_i is linked to investments in counter-espionage aimed at preventing the opponent from spying on her primitive information, then a higher t_i , by reflecting a smaller investment in such defensive measures, naturally comes with a lower cost.

Given any pair of cognitive profiles $\rho = (\rho_1, \rho_2)$, with $\rho_i = (l_i, t_i)$, $i = 1, 2$, let

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

denote the endogenous precisions of the total noise $\eta_i \equiv \phi_i + \gamma_j$ and $\eta_j \equiv \phi_j + \gamma_i$ in the two players' secondary signals. Fixing $\rho = (\rho_1, \rho_2)$, 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 $s_j = (s_j^P, s_j^S)$, with probability one an action $a_j(s_j) = m_j s_j^P + (1 - m_j) s_j^S$ that is a convex combination of s_j^P and s_j^S , her best response is to follow a stage-2 strategy that, for each $s_i = (s_i^P, s_i^S)$, also selects with probability one an action $a_i(s_i) = m_i s_i^P + (1 - m_i) s_i^S$ that is a convex combination of s_i^P and s_i^S , with

$$m_i = \frac{1 + r_i^\rho(1 + \beta) - 2\beta r_i^\rho m_j}{1 + 2r_i^\rho}. \quad (10)$$

When actions in the downstream game are 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). This property naturally reflects the desire for player i to align her action with player j 's action. The opposite is true in case of strategic substitutes ($\beta < 0$).

Furthermore, given any pair of cognitive profiles $\rho = (\rho_1, \rho_2)$, there exists a unique linear continuation equilibrium $\sigma^\rho = (\sigma_1^\rho, \sigma_2^\rho)$ for the stage-2 game and is such that, for each $s_i = (s_i^P, s_i^S)$, σ_i^ρ selects with certainty the equilibrium action $a_i^\rho(s_i) = m_i^\rho s_i^P + (1 - m_i^\rho) s_i^S$ with

$$m_i^\rho = \frac{1 + 2r_j^\rho + r_i^\rho(1 - \beta) + 2r_i^\rho r_j^\rho(1 - \beta^2)}{1 + 2(r_i^\rho + r_j^\rho) + 4r_i^\rho r_j^\rho(1 - \beta^2)}. \quad (11)$$

One can also verify that, when the two players are expected to engage in cognition $\rho = (\rho_i, \rho_j)$ and, instead, player i selects cognition $\rho'_i = (l'_i, t'_i)$, in the stage-2 game, for any $s_i = (s_i^P, s_i^S)$, player i 's best response to player j following the equilibrium strategy σ_j^ρ consists in selecting with certainty the action $a_i^{\rho'_i; \rho}(s_i) = m_i^{\rho'_i; \rho} s_i^P + (1 - m_i^{\rho'_i; \rho}) s_i^S$, 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 action to his primary signal and where

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

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

4.1 Espionage

We now specialize the noisy-communication model above to the case in which cognition is purely self-directed: players “spy” on each other, that is, take actions to gather (noisy) information about the opponents' beliefs and information (formally captured by their primitive signals). We abstract from the possibility of counter-espionage, which we address separately in the next subsection.

We capture such a situation as follows. We assume that the precision t_j of the component γ_j in the noise $\eta_i \equiv \phi_i + \gamma_j$ in player i 's secondary signal $s_i^S = s_j^P + \eta_i$ is infinity so that γ_j is identically equal to zero in which case the noise η_i coincides with ϕ_i . Abusing notation, we then drop ϕ_i and let player i 's cognition ρ_i coincides with the precision of the noise η_i in player i 's secondary signal. That is, each player directly controls the precision of her secondary signal. As anticipated above, such a secondary signal can be interpreted as the outcome 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 cognitive 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 a noise is exogenous, and Calvo-Armengol et al., 2015, and Sethi and Yildiz, 2016, 2018, for models in which players choose the precision of the information they receive from other players).

We then have the following result:

Proposition 6 (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 $\beta(\hat{\rho}_i - \rho_i)(\hat{\rho}_j - \rho_j) \leq 0$.*

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.²⁸ When $\beta > 0$, this induces player j to rely

²⁸To see this, note that the formula in (11), when specialized to the case of purely self-directed cognition (that is, to the model of espionage under consideration here), is equal to

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

It is then easy to verify that m_i^ρ is decreasing in ρ_i , which means that the sensitivity $1 - m_i^\rho$ of player i 's stage-2 action to her secondary signal s_i^S is increasing in ρ_i .

more on her primary signal—player j likes being spied.²⁹ In turn, 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: this is because player i can now learn player j 's information without ending up aligning her action much with player j 's. Thus UEC holds in these games, no matter whether the downstream actions are strategic complements or substitutes.

Next, consider ID. 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 relies more on her secondary signal and less on her primary one. When $\beta > 0$, this reduces player i 's incentives to spy, whereas the opposite is true when $\beta < 0$.³⁰

4.2 Information sharing and counter-espionage

We now consider the opposite polar case of the general noisy-communication model above in which cognition is purely manipulative. Formally, we assume that the precisions l_i of the terms ϕ_i in the agents' secondary signals are equal to infinity, in which case the terms ϕ_i are identically equal to zero. The noise η_i in player i 's secondary signal then coincides with γ_j and is thus entirely controlled by player j (purely manipulative cognition). As anticipated above, this situation is meant to capture instances in which players take actions to facilitate their opponents' ability to understand their views of the game or, alternatively, invest in *counter-espionage* to make it difficult for their opponents to spy on them.

For any $\rho'_i = (t'_i, l'_i)$ and $\rho = (\rho_i, \rho_j)$, with $\rho_i = (t_i, l_i)$ and $\rho_j = (t_j, l_j)$, we then have that the precisions of the noise η_i and η_j in the secondary signals are equal to

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

Abusing notation again, we then let each player's cognition coincide with the precision of the noise in the opponent's secondary signal. We interpret a larger ρ_i as a larger effort to make oneself understood or, equivalently, a smaller ρ_i as a higher investment in counter-espionage, that is, in activities that make it difficult for player j to spy on i 's primitive information.

We then have the following result:

Proposition 7 (counter-espionage). *Let $\rho = (\rho_i, \rho_j)$ and $\hat{\rho} = (\hat{\rho}_i, \hat{\rho}_j)$. UEC holds for such*

²⁹To see this, note that the formula in (10), when specialized to the case of purely self-directed cognition (that is, to the model of espionage under consideration here), and applied to player j , is equal to

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

It is immediate that m_j is decreasing in m_i and hence increasing in $1 - m_i$.

³⁰One 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. These conditions, however, are not particularly illuminative and hence we do not discuss them here.

cognitive profiles if $\beta > 0$, and does not hold if $\beta < 0$. Irrespectively 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 case of espionage considered in Subsection 4.1, whether or not UEC holds in the counter-espionage game 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 player i 's secondary signal $s_i^S = s_j^P + \gamma_j$. When player j expects player i to invest less in counter-espionage (equivalently, to pass on a more precise secondary 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.³¹ When actions are complements ($\beta > 0$), in turn this increases the incentives for player i to reduce her investment in counter-espionage (equivalently, to send player j a more precise secondary signal so as to better coordinate with her), whereas the opposite is true when downstream 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 reciprocate by sending player j a more precise secondary signal are smaller, implying that this game satisfies a negative form of increasing differences: players share less information with the rivals when they expect the latter to share more information with them.

Other things equal, the above properties imply that, in the case of strategic complements, counter-espionage is likely to give rise to equilibria with asymmetric cognitive postures.

5 Other forms of cognition

In this section we discuss the role that EC, and its decomposition into UEC and ID, play in other settings. Because of space constraints, the exposition in this section is less formal. We report the key insights but refer the reader to the Online Supplement for the details.

5.1 Self-direct cognition: Noisy information about exogenous payoff states

In most of the literature on information acquisition in games, cognition is self-directed and takes the form of players receiving additive signals about an unknown exogenous payoff state, with the noise in each player's signal determined by the player's own cognition (see, among others, Hellwig and Veldkamp 2009, Myatt and Wallace, 2012, Colombo et al, 2014, Szkup and Trevino, 2015, and Pavan,

³¹To see this, note that the formula in (11), when specialized to the case of purely manipulative cognition (that is, to the model of counter-espionage under consideration here), and applied to player j ,

is equal to

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

It is then easy to see that the sensitivity 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 .

2016). In the Online Supplement, we consider the role that EC plays in such games. Specifically, we assume that payoffs take the same form as in Section 4 but each player receives one perfectly public and one perfectly private signal, with the precision of the public signal exogenous and with that of the private signal determined endogenously by the player’s cognition.

We show that, no matter whether actions in the downstream game are strategic complements or substitutes, UEC always holds in these games. ID, instead, holds for strategic complements but not for substitutes.

The literature on noisy information acquisition in linear-quadratic-Gaussian games such as those considered in the Online Supplement has pointed out that such games may or may not feature multiple equilibria, depending on the structure of the information cost C . For example, when actions are complements, Hellwig and Veldkamp (2009) showed that such games typically admit multiple (symmetric) equilibria when players have access to a finite set of information sources, whereas Myatt and Wallace (2012) showed 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. Our results in the Online Supplement contribute a different angle. First, they shed light on the players’ incentives to conform to their opponents’ expectations, for *arbitrary cognitive profiles*, which need not be symmetric across the players. More importantly, they help identify the underlying forces contributing to equilibrium multiplicity, 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’ cognitive postures (ID). The decomposition provides a deeper understanding of the nature of the equilibrium determinacy in these games.

5.2 Manipulative cognition: framing and generalized career concerns

In many situations of interest, some players use “frames” (broadly interpreted as the design of contextual purchasing experiences or other manipulative devices aimed at inducing other agents to act favorably to them). In the Online Supplement, we consider a simplified persuasion game in which a Sender (the persuader) manipulates a Receiver’s recollection of information relevant for a decision.³² The analysis is motivated by the design of contextual purchasing experiences, in the spirit of Salant and Siegel (2018).

What distinguishes these situations from those examined in the Bayesian persuasion literature (see, Bergemann and Morris, 2019, and Kamenica, 2019, for overviews) is the inability of the Sender to commit to her choice of a frame (i.e., to her information structure). We show that the Sender’s incentives to use manipulative frames are stronger when she is expected to invest more into manipulative framing. This is because, the more the Receiver expects the Sender to engage in manipulative framing, the more she interprets the lack of recollection of information favorable to the Sender as a signal of the state being unfavorable to the Sender. But then the stronger the incentives for the

³²See also Schwartzstein and Sunderam (2021) for an alternative model of manipulative persuasion in which the Sender influences the receivers by proposing models that organize past data to make predictions. Related is also Gibbons et al. (2021). That paper studies the effects of frames (interpreted as simplified descriptions of complex strategic situations) on the collective performance of a group of agents.

Sender to engage into manipulative framing to avoid such a situation. As a result, these games naturally feature UEC for the Sender.

When the Receiver can invest to increase her ability to recollect information previously received (and hence limit the influence of the Sender’s manipulation), we also identify conditions under which ID holds (i.e., the Sender’s incentives to engage in manipulative framing are stronger when the Receiver invests more in defensive memory management), and show that such conditions naturally relate to the (sub)modularity of the Receiver’s memory technology in her memory investment and in the Sender’s framing effort.

Manipulative framing is an instance of “signal jamming,” akin to those studied in Holmström (1999)’s celebrated model of career concerns.³³ In the Online Supplement, we consider a generalization of this model that encompasses most of the cases considered in the signal-jamming literature. We show that when the worker’s talent and effort are substitutes as in Holmström (1999) additive model, EC never obtains. When, instead, they are complements, as in Dewatripont et al (1999)’s multiplicative model, EC typically obtains. Consistently with the results in Proposition 1, the equilibrium is thus unique in Holmström (1999), whereas multiple equilibria are possible in Dewatripont et al. (1999).³⁴ We also show that, when this is the case, the worker is typically better off in the low-effort equilibrium (a form of cognitive trap). Whether or not the above conclusions are robust to the possibility that the labor market also invests in information acquisition then depends on the specific assumptions one makes on the output and information-acquisition technologies, as we discuss in the Online Supplement.³⁵

EC plays an important role also in many other environments in which cognition is manipulative. For example, consider a trading environment in which a seller invests in information acquisition and then decides when to disclose hard information to the buyer to convince her to trade (see, e.g., Pavan and Tirole, 2022b). EC arises naturally in this situation. If the seller is expected by the buyer to exert substantial effort, the price when the seller does not disclose hard evidence that the buyer’s value is high is low, in which case it is particularly profitable for the seller to invest to convince the buyer that her value is high when this is indeed the case (UEC). In case of multiplicity, the seller is again typically better off in a low-effort equilibrium. So expectation traps naturally arise in this context as well.

³³Signal jamming has also received attention in the Industrial Organization literature. It occurs, for example, 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. See Fudenberg and Tirole (1986).

³⁴See also Hörner and Lambert (2019) for a more recent analysis of these games.

³⁵Another class of signal-jamming games giving rise to EC is intra-personal memory management. In these games, a player receives information that she may try to remember, or repress, to influence the behavior of her future selves. Both repression and cognitive discipline are typically dictated by the expectations of future selves, with distinct welfare implications. These 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.

5.3 Endogenous depth of reasoning

The analysis in all the situations discussed above assumes that the players are fully rational, and that cognition takes the form of learning about payoffs and/or about other players' beliefs. In the Online Supplement, we also consider an alternative class of cognitive games in which payoffs are common knowledge, but where players are boundedly rational and 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. The earlier literature (see, e.g., Crawford, Costa-Gomes, and Iriberry, 2013, for a detailed overview) assumes that a player's depth of reasoning is exogenous. Alaoui and Penta (2016, 2017, 2018) are the first to endogenize the depth of reasoning in the level- k model. Our analysis differs from theirs in 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 about her own depth of reasoning.

A key difference relative to the settings considered above is that, in this model, a player understands that, by increasing her cognition, she may end up with a lower payoff. 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 mixed action that is farther away from her true best response than the one identified by computing a smaller number of iterated best responses. The model also shares a few notable features with the sparsity model: (i) a player who goes significantly deeper into the understanding of the game than her opponent does not gain any advantage in predicting her opponent's behavior vis-a-vis a player who goes slightly deeper; (ii) 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.

We show that EC plays a key role also in this model. In particular, we consider a simplified version of the celebrated Arad and Rubinstein (2012)'s 11-20 game. This game, which is intended for experimental work, captures, in a stark and simplified manner, some of the forces that arise in strategic situations where players benefit from matching, or undercutting by little, the rivals' actions (e.g., oligopoly games with imperfect substitutable products where firms compete in prices).

We show that this game exhibits a negative form of EC: a player's incentives to go deeper in the understanding of the process of iterated best responses are lower the higher the depth of reasoning expected from her by the opponent (UEC) and the higher the opponent's depth of reasoning. Arguments similar to those establishing Proposition 1 above then imply that, when the cognitive costs are strictly increasing, this model never features multiple (pure-strategy) equilibria.

6 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 players or, more generally, invest to learn other players' view of the game, but also take manipulative actions aimed at

facilitating (alternatively, obstructing) other players' understanding of their own view of the game, as in the case of counter-espionage.

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 two-player-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, and information acquisition about other players' beliefs as in the case of espionage and counter-espionage.

While the concept of expectation conformity accommodates for the possibility that a player's cognition is partly observable by the other players and dynamic, throughout the analysis, we focus on situations in which the cognitive investments are not observable by other players and essentially static. In future work, it would be interesting to develop a few applications in which cognition is publicly observable and 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 normative aspects related to the selection of the cognitive postures. 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(4), 1297–1333.
- Alaoui, L. and A. Penta (2022), “Cost-Benefit Analysis in Reasoning,” *Journal of Political Economy*, 130(4) 881-925.
- Alaoui, L., K.A. Janezic, and A. Penta (2020), “Reasoning about Others’ Reasoning,” *Journal of Economic Theory*, 189, 105091.
- 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 K. Sastry (2021), “Inattentive Economies,” mimeo MIT.
- 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., J. Davis, and N. Gondhi (2021), “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(2): 429–462.
- Bénabou, R., and J. Tirole (2002) “Self-Confidence and Personal Motivation,” *Quarterly Journal of Economics*, 117(3): 871–915.
- Bénabou, R. and J. Tirole (2006) “Belief in a Just World and Redistributive Politics,” *Quarterly Journal of Economics*, 121(2): 699-746.
- 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(1), 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, (2022), “Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy,” *Journal of Political Economy*, 130(6), 1676-1715.
- Carroll, G. (2019), “Information Games and Robust Trading Mechanisms,” mimeo University of Toronto.
- Che, Y.K., and N. Kartik (2009) “Opinions as Incentives,” *Journal of Political Economy*, 117(5): 815–860.

- Colombo, L., G. Femminis, and A. Pavan (2014) "Information Acquisition and Welfare," *Review of Economic Studies*, 81(4): 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., F. Khalil, and J.C. Rochet (1998a) "Contracts and Productive Information Gathering." *Games and Economic Behavior*, 25(2): 174–193.
- Crémer, J., F. Khalil, and J.C. Rochet (1998b) "Strategic Information Gathering before a Contract is Offered." *Journal of Economic Theory*, 81(1): 163–200.
- Crawford, V.P., M.A. 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(1): 339–354.
- Dang, T.V., G. Gorton, B. Holmström, and G. Ordonez (2017) "Banks as Secret Keepers," *American Economic Review*, 107(4): 1005–1029.
- Dessi, R. (2008) "Collective Memory, Cultural Transmission and Investments," *American Economic Review*, 98(1): 534–560.
- Denti, T. (2022), "Unrestricted Information Acquisition," *Theoretical Economics*, forthcoming.
- Dewatripont, M., I. Jewitt, and J. Tirole (1999) "The Economics of Career Concerns, Part I: Comparing Information Structures," *Review of Economic Studies*, 66 (1): 183–198.
- Dewatripont, M., and J. Tirole (2005) "Modes of Communication," *Journal of Political Economy*, 113(6): 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*, 1(1): 60–79.
- Gibbons, R., M. LiCalzi, and M. Warglien (2021), "What Situation Is This? Shared Frames and Collective Performance," *Strategy Science* 6(2):124-140.
- Gottlieb, D. (2014) "Imperfect Memory and Choice under Risk," *Games and Economic Behavior*, 85: 127–185.
- Gottlieb, D. (2022) "Will You Never Learn? Self-Deception and Biases in Information Processing," mimeo, Olin Business School, Washington University in St. Louis.
- Hébert, B. and J. La'O (2021) "Information Acquisition, Efficiency, and Non-Fundamental Volatility," NBER Working Paper No. 26771.
- Hébert, B. and M. Woodford (2021), "Neighborhood-Based Information Costs," *American Economic Review*, 111(10): 3225–3255.

- Heller, Y. and E. Winter (2016), “Rule Rationality,” *International Economic Review*, 57(3), 997-1026.
- 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.
- Hörner, J. and N.S. Lambert, (2021), “Motivational Ratings,” *Review of Economic Studies*, 88(4): 1892–1935.
- Kamenica, E., (2019), “Bayesian Persuasion and Information Design,” *Annual Review of Economics*, 11: 249–272.
- 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(3): 343–357.
- Liang, A., and X. Mu (2020) “Complementary Information and Learning Traps,” *Quarterly Journal of Economics*, 135(1): 389–448.
- Maćkowiak, B., F. Matějka, and M. Wiederholt, (2022) “Rational Inattention: A Review,” *Journal of Economic Literature*, forthcoming.
- Maćkowiak, B., and M. Wiederholt (2009) “Optimal Sticky Prices under Rational Inattention,” *American Economic Review*, 99(3): 769–803.
- Matějka, F., and A. McKay (2015) “Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model,” *American Economic Review*, 105(1): 272-298.
- Morris, S., and H. S. Shin (2002) “Social Value of Public Information,” *American Economic Review*, 92(5): 1521–1534.
- Morris, S., and M. Yang (2022) “Coordination and Continuous Stochastic Choice,” *Review of Economic Studies*, 89(5): 2687-2722.
- Mullainathan, S. (2002) “Thinking through Categories”, mimeo, MIT.
- Mullainathan, S., J. Schwartzstein, and A. Shleifer (2008) “Coarse Thinking and Persuasion,” *Quarterly Journal of Economics*, 123(2): 577–619.
- Myatt, D. P., and C. Wallace (2012) “Endogenous Information Acquisition in Coordination Games,” *Review of Economic Studies*, 79(1): 340–374.
- Pavan, A. (2016) “Attention, Coordination, and Bounded Recall,” mimeo, Northwestern University.
- Pavan, A., and J. Tirole, (2022a) “Exposure to the Unexpected, Duty of Disclosure and Contract Design,” mimeo, Northwestern University and Toulouse School of Economics.
- Pavan, A., and J. Tirole (2022b) “Knowing your Lemon before you Dump It,” mimeo, Northwestern University and 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): 315-346.

- Schwartzstein, J., and A. Sunderam (2021), “Using Models to Persuade,” *American Economic Review*, 111 (1): 276-323.
- Sims, C. A. (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, Columbia University and MIT.
- Steiner J., C. Stewart, and F. Matějka, (2017) “Rational Inattention Dynamics: Inertia and Delay in Decision-making,” *Econometrica*, 85(2), 521–553.
- Stigler, G. J. (1961) “The Economics of Information,” *Journal of Political Economy*, 69(3): 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 3. 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. We then establish the properties in the proposition.

First note that the players' payoffs can be rewritten as

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

Hereafter, we disregard the terms in the curly bracket because they have no impact on individual best responses and cancel out when computing the differential payoffs in the definition of UEC and ID. Using the fact that $g(\omega) = (1 + \sum_{k=1}^K \omega^k)/(1 - \beta)$, we then proceed as if the agents' payoffs were given by

$$\hat{u}_i(a_i, a_{-i}, \omega) = 2a_i \left(1 + \sum_{k=1}^K \omega^k + \beta\bar{a}_{-i}\right) - a_i^2. \quad (12)$$

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, $(\omega^k)_{k=1}^K$ as the components to the gross return that the players learn about, $\beta\bar{a}_{-i}$ as an investment spillover (positive for $\beta > 0$ and negative for $\beta < 0$), and a_i^2 as a quadratic adjustment cost.

As mentioned in the main text, to further simplify the derivations, we assume that there are only two players, so that $n = 2$.

Downstream best responses, equilibrium actions, and equilibrium payoffs.

Fix the cognitive profile $\rho = (\rho_1, \rho_2)$, with $\rho_1 \leq \rho_2$, and note that, given any ω , and $s_i = (\omega^1, \dots, \omega^{\rho_i})$, the two players' best responses must satisfy the optimality conditions

$$a_i = 1 + \bar{\omega}^{\rho_i} + \beta \mathbb{E}[a_j | \rho_i, s_i]$$

for $i, j = 1, 2, j \neq i$, where

$$\bar{\omega}^{\rho_i} \equiv \sum_{k=1}^{\rho_i} \omega^k.$$

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. Hereafter, we denote the stage-2 equilibrium actions by $a_i^\rho(s_i)$.

That player 2 goes deeper in the exploration of the state than player 1 implies that, for any $s_1 = (\omega^1, \dots, \omega^{\rho_1})$ and $s_2 = (\omega^1, \dots, \omega^{\rho_2}) = (s_1, \omega^{\rho_1+1}, \dots, \omega^{\rho_2})$, $\mathbb{E}[a_1^\rho(s_1) | \rho_2, s_2] = a_1^\rho(s_1)$. Player's 2 stage-2 equilibrium action must thus satisfy the optimality condition

$$a_2^\rho(s_2) = 1 + \bar{\omega}^{\rho_2} + \beta a_1^\rho(s_1)$$

for all $s_2 = (\omega^1, \dots, \omega^{\rho_2})$. In turn, this implies that, for any $s_1 = (\omega^1, \dots, \omega^{\rho_1})$, player 1's stage-2 equilibrium action must satisfy

$$a_1^\rho(s_1) = 1 + \bar{\omega}^{\rho_1} + \beta [1 + \bar{\omega}^{\rho_1} + \beta a_1^\rho(s_1)].$$

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

$$a_1^\rho(s_1) = \frac{1 + \bar{\omega}^{\rho_1}}{1 - \beta}$$

and

$$a_2^\rho(s_2) = a_1^\rho(s_1) + \bar{\omega}^{\rho_2} - \bar{\omega}^{\rho_1}.$$

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

$$\mathbb{E} [\hat{u}_1(a_1^\rho(s_1), a_2^\rho(s_2), \omega) | \rho_1, s_1] = 2a_1^\rho(s_1) (1 + \bar{\omega}^{\rho_1} + \beta a_1^\rho(s_1)) - (a_1^\rho(s_1))^2 = \left(\frac{1 + \bar{\omega}^{\rho_1}}{1 - \beta} \right)^2$$

and

$$\mathbb{E} [\hat{u}_2(a_1^\rho(s_1), a_2^\rho(s_2), \omega) | \rho_2, s_2] = 2a_2^\rho(s_2) (1 + \bar{\omega}^{\rho_2} + \beta a_1^\rho(s_1)) - (a_2^\rho(s_2))^2 = \left(\frac{1 + \bar{\omega}^{\rho_2} - \beta(\bar{\omega}^{\rho_2} - \bar{\omega}^{\rho_1})}{1 - \beta} \right)^2.$$

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 in a fictitious environment in which (a) all dimensions ω_k , $k > \rho_1$, are equal to zero and such an 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 in a fictitious environment in which (a) all dimensions $k > \rho_2$ are equal to zero and such an 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 "do not exist", as claimed in the main text.

Player 1's deviations.

Now suppose that the two players are expected to engage in cognition $\rho = (\rho_1, \rho_2)$, with $\rho_1 < \rho_2$, and player 1 deviates and selects cognition ρ'_1 with $\rho_1 < \rho'_1 \leq \rho_2$. To shorten the formulas, with an abuse of notation, let

$$s_1 = (\omega^1, \dots, \omega^{\rho_1})$$

and

$$s'_1 = (\omega^1, \dots, \omega^{\rho_1}, \omega^{\rho_1+1}, \dots, \omega^{\rho'_1}) = (s_1, \omega^{\rho_1+1}, \dots, \omega^{\rho'_1}).$$

Player 1's optimal stage-2 strategy then selects, for any s'_1 , with probability one the action

$$\begin{aligned} a_1^{\rho'_1; \rho}(s'_1) &= 1 + \bar{\omega}^{\rho'_1} + \beta \mathbb{E} [a_2^\rho(s_2) | \rho'_1, s'_1] &= 1 + \bar{\omega}^{\rho'_1} + \beta \mathbb{E} [a_2^\rho(s_2) | \rho'_1, s'_1] \\ &= 1 + \bar{\omega}^{\rho'_1} + \beta (a_1^\rho(s_1) + \bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_1}) \\ &= a_1^\rho(s_1) + (1 + \beta)(\bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_1}). \end{aligned}$$

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

$$\mathbb{E} \left[\hat{u}_1(a_1^{\rho'_1; \rho}(s'_1), a_2^\rho(s_2), \omega) | \rho'_1, s'_1 \right] = 2a_1^{\rho'_1; \rho}(s'_1) \left[1 + \bar{\omega}^{\rho'_1} + \beta \left(a_1^\rho(s_1) + \bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_1} \right) \right] - \left(a_1^{\rho'_1; \rho}(s'_1) \right)^2$$

which simplifies to

$$\begin{aligned} \mathbb{E} \left[\hat{u}_1(a_1^{\rho'_1; \rho}(s'_1), a_2^\rho(s_2), \omega) | \rho'_1, s'_1 \right] &= 2a_1^\rho(s_1) (1 + \bar{\omega}^{\rho_1} + \beta a_1^\rho(s_1)) - (a_1^\rho(s_1))^2 \\ &+ 2a_1^\rho(s_1)(1 + \beta)(\bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_1}) + 2(1 + \beta)(\bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_1}) (1 + \bar{\omega}^{\rho_1} + \beta a_1^\rho(s_1)) \\ &+ 2(1 + \beta)^2(\bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_1})^2 - (1 + \beta)^2(\bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_1})^2 - 2a_1^\rho(s_1)(1 + \beta)(\bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_1}). \end{aligned}$$

Integrating over states, we then have that player 1's 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 $\rho'_1 > \rho_2$. Continue to let

$$s_1 = (\omega^1, \dots, \omega^{\rho_1})$$

and

$$s'_1 = (\omega^1, \dots, \omega^{\rho_1}, \omega^{\rho_1+1}, \dots, \omega^{\rho'_1}) = (s_1, \omega^{\rho_1+1}, \dots, \omega^{\rho'_1}).$$

Then, for any s'_1 , player 1's optimal stage-2 strategy consists in choosing with probability one the action

$$\begin{aligned} a_1^{\rho'_1; \rho}(s'_1) &= 1 + \bar{\omega}^{\rho'_1} + \beta \mathbb{E} [a_2^\rho(s_2) | \rho'_1, s'_1] \\ &= 1 + \bar{\omega}^{\rho_1} + \bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_1} + \beta (a_1^\rho(s_1) + \bar{\omega}^{\rho_2} - \bar{\omega}^{\rho_1}) \\ &= a_1^\rho(s_1) + (1 + \beta)(\bar{\omega}^{\rho_2} - \bar{\omega}^{\rho_1}) + \bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_2} \\ &= a_2^\rho(s_2) + \beta(\bar{\omega}^{\rho_2} - \bar{\omega}^{\rho_1}) + \bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_2}. \end{aligned}$$

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

$$\mathbb{E} \left[\hat{u}_1(a_1^{\rho'_1; \rho}(s'_1), a_2^\rho(s_2), \omega) | \rho'_1, s'_1 \right] = 2a_1^{\rho'_1; \rho}(s'_1) \left(1 + \bar{\omega}^{\rho'_1} + \beta a_2^\rho(s_2) \right) - \left(a_1^{\rho'_1; \rho}(s'_1) \right)^2$$

which simplifies to

$$\begin{aligned} \mathbb{E} \left[\hat{u}_1(a_1^{\rho'_1; \rho}(s'_1), a_2^\rho(s_2), \omega) | \rho'_1, s'_1 \right] &= 2a_2^\rho(s_2) (1 + \bar{\omega}^{\rho_2} + \beta a_1^\rho(s_1)) - (a_2^\rho(s_2))^2 \\ &+ 2\beta(\bar{\omega}^{\rho_2} - \bar{\omega}^{\rho_1}) (1 + \bar{\omega}^{\rho_1} + \beta a_1^\rho(s_1)) + 2\beta(1 + \beta)(\bar{\omega}^{\rho_2} - \bar{\omega}^{\rho_1})^2 + 2\beta(\bar{\omega}^{\rho_2} - \bar{\omega}^{\rho_1})(\bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_2}) \\ &+ 2(\bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_2}) [1 + \bar{\omega}^{\rho_2} + \beta (a_1^\rho(s_1) + \bar{\omega}^{\rho_2} - \bar{\omega}^{\rho_1})] + 2(\bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_2})^2 \\ &- \beta^2(\bar{\omega}^{\rho_2} - \bar{\omega}^{\rho_1})^2 - (\bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_2})^2 - 2\beta(\bar{\omega}^{\rho_2} - \bar{\omega}^{\rho_1})(\bar{\omega}^{\rho'_1} - \bar{\omega}^{\rho_2}). \end{aligned}$$

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

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

Finally, suppose that the two players are expected to choose cognition $\rho = (\rho_1, \rho_2)$, with $\rho_1 \leq \rho_2$, and player 1 deviates to some lower cognition $\rho'_1 < \rho_1$. Let

$$s'_1 = (\omega^1, \dots, \omega^{\rho'_1})$$

and

$$s_1 = (\omega^1, \dots, \omega^{\rho'_1}, \omega^{\rho'_1+1}, \dots, \omega^{\rho_1}) = (s'_1, \omega^{\rho'_1+1}, \dots, \omega^{\rho_1}).$$

In this case, player 1's optimal action is equal to

$$\begin{aligned} a_1^{\rho'_1; \rho}(s'_1) &= 1 + \bar{\omega}^{\rho'_1} + \beta \mathbb{E}[a_2^\rho(s_2) | \rho'_1, s'_1] \\ &= 1 + \bar{\omega}^{\rho'_1} + \beta \mathbb{E}[a_1^\rho(s_1) + \bar{\omega}^{\rho_2} - \bar{\omega}^{\rho_1} | \rho'_1, s'_1] \\ &= 1 + \bar{\omega}^{\rho'_1} + \beta \frac{1 + \bar{\omega}^{\rho_1}}{1 - \beta} \\ &= a_1^{\rho'_1, \rho_2}(s'_1). \end{aligned}$$

It is evident that player 1's ex-ante expected payoff is then equal to

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

Player 2's deviations.

Next, consider deviations by the player who is expected to be ahead in the exploration of the state. Namely, 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 $\rho'_2 \geq \rho_1$. Let

$$s_1 = (\omega^1, \dots, \omega^{\rho_1})$$

and

$$s'_2 = (\omega^1, \dots, \omega^{\rho_1}, \omega^{\rho_1+1}, \dots, \omega^{\rho'_2}) = (s_1, \omega^{\rho_1+1}, \dots, \omega^{\rho'_2}).$$

Player 2's optimal stage-2 action is then equal to

$$a_2^{\rho'_2; \rho}(s'_2) = 1 + \bar{\omega}^{\rho'_2} + \beta \mathbb{E}[a_1^\rho(s_1) | \rho'_2, s'_2] = 1 + \bar{\omega}^{\rho'_2} + \beta a_1^\rho(s_1) = a_2^{\rho_1, \rho'_2}(s'_2).$$

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 a deviation is thus equal to

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

Finally, consider the payoff that player 2 obtains by deviating to cognition $\rho'_2 < \rho_1$. Let

$$s'_2 = (\omega^1, \dots, \omega^{\rho'_2})$$

and

$$s_1 = (\omega^1, \dots, \omega^{\rho'_2}, \omega^{\rho'_2+1}, \dots, \omega^{\rho_1}).$$

Player 2's optimal stage-2 action is then equal to

$$a_2^{\rho'_2; \rho}(s'_2) = 1 + \bar{\omega}^{\rho'_2} + \beta \mathbb{E}[a_1^\rho(s_1) | \rho'_2, s'_2] = 1 + \bar{\omega}^{\rho'_2} + \beta \frac{1 + \bar{\omega}^{\rho'_2}}{1 - \beta} = \frac{1 + \bar{\omega}^{\rho'_2}}{1 - \beta} = a_2^{\rho_1; \rho'_2}(s'_2).$$

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

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

and

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

Hence UEC holds strictly for player 1 and weakly (i.e., as an equality) for player 2.

Next, consider ID. 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 4. First, we prove that all equilibria are symmetric. To see this, use the derivations in the proof of Proposition 3 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 cognitive choices. Thus assume an equilibrium exists in which $\rho_1 < \rho_2$. The arguments in the proof of Proposition 3 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 3 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 3 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 last term 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 $k^* \geq \bar{k} \equiv \min \left\{ k \mid \sigma_k^2 \leq c_k \right\}$.

Next, suppose that player 1 deviates to cognition $\rho'_1 \leq k^* - 1$. The derivations in the proof of Proposition 3 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 last term 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, for any $k^* \in [\underline{k}, \bar{k}(\beta)] \cap \mathbb{N}$, a symmetric equilibrium with cognition k^* exists, and there is no symmetric equilibrium with cognition $k^* < \underline{k}$ or $k^* > \bar{k}(\beta)$.

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 $\psi(a_{-i}, \omega) = (1 - \beta)g(\omega)^2 + \beta\bar{a}_{-i}^2$). That equilibria are Pareto ranked then follows from the fact that, in the symmetric equilibrium with depth of knowledge k^* , the equilibrium payoffs 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 all the equilibria of the game in which cognition is ordered also also equilibria in the game in which cognition is unordered and that, under Condition (5), the converse is also true follows from the arguments in the main text. Q.E.D.

Proof of Proposition 5. From the arguments in the proof of Proposition 4, 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,

$$\frac{\sigma_{k^*+1}^2}{c_{k^*+1}} \leq 1 \leq \frac{\sigma_{k^*}^2}{(1 - \beta)^2 c_{k^*}}. \quad (13)$$

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 4).

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 cognitive level $\rho'_2 \geq \rho_1$ is equal to

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

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,

$$\frac{\sigma_{\rho_2+1}^2}{c_{\rho_2+1}} \leq 1 \leq \frac{\sigma_{\rho_2}^2}{c_{\rho_2}}. \quad (14)$$

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 3, 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

Thus, consider player 1 (the one who selects the lowest cognitive level). As shown in the proof of Proposition 3, 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

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

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,

$$\frac{\sigma_{\rho_1+1}^2 (1 + \beta)^2}{c_{\rho_1+1}} \leq 1 \leq \frac{\sigma_{\rho_1}^2}{(1 - \beta)^2 c_{\rho_1}}. \quad (15)$$

The two inequalities in (15), 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 (14), Condition (15) also implies that player 1 prefers ρ_1 to any $\rho'_1 > \rho_2$.

We thus conclude that Conditions (14) and (15), which are equivalent to Conditions (8) and (7) in the main text, are both necessary and sufficient for existence of an 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 $[\underline{k}(\beta), \bar{k}] \cap \mathbb{N}$ then follows from the above results along with the fact that $\rho_1 < \rho_2$.

Next, 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

³⁶Indeed, from the proof of Proposition 3, 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 $\rho_2 \geq \rho_1$.

are multiple solutions ρ_2 to the double inequality in (14), is invariant in ρ_2 . It is also easy to see that the sum of the two players' payoffs is maximal under the equilibrium featuring the lowest cognition for player 1 (the one learning the smallest number of dimensions).

That all the equilibria of the game in which cognition is ordered are also equilibria in the game in which cognition is unordered and, under Condition (9), that the converse is also true follows from the arguments in the main text. Q.E.D.

Derivation of the conditions at the beginning of Section 4

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

$$\mathbb{E} [\omega | s_i; \rho'_i, \rho_j] = \frac{1}{1 + h_i^S(\rho'_i, \rho_j)} s_i^P + \frac{h_i^S(\rho'_i, \rho_j)}{1 + h_i^S(\rho'_i, \rho_j)} s_i^S$$

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

$$h_i^S(\rho'_i, \rho_j) \equiv [\text{var}(\varepsilon_j + \gamma_j + \phi_i) | \rho'_i, \rho_j]^{-1} = \frac{r_i^{\rho'_i, \rho_j}}{1 + r_i^{\rho'_i, \rho_j}}$$

is the total precision of player i 's secondary signal, with

$$r_i^{\rho'_i, \rho_j} \equiv [\text{var}(\gamma_j + \phi_i) | \rho'_i, \rho_j]^{-1} = \left[\frac{1}{t_j} + \frac{1}{l'_i} \right]^{-1} = \frac{t_j l'_i}{t_j + l'_i}.$$

Similarly,

$$\mathbb{E} [s_j^S | s_i; \rho'_i, \rho_j] = s_i^P$$

and

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

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

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

for some scalar m_j . Given ρ'_i , when player i expects player j to engage in cognition ρ_j , then, for any $s_i = (s_i^P, s_i^S)$, player i 's best response in the stage-2 game consists in choosing with certainty the

action

$$\begin{aligned}
a_i &= (1 - \beta)\mathbb{E}[\omega|s_i; \rho'_i, \rho_j] + \beta\mathbb{E}[a_j(s_j)|s_i; \rho'_i, \rho_j] \\
&= (1 - \beta)\frac{1+r_i^{\rho'_i, \rho_j}}{1+2r_i^{\rho'_i, \rho_j}}s_i^P + (1 - \beta)\frac{r_i^{\rho'_i, \rho_j}}{1+2r_i^{\rho'_i, \rho_j}}s_i^S \\
&\quad + \beta m_j \left[\frac{1}{1+2r_i^{\rho'_i, \rho_j}}s_i^P + \frac{2r_i^{\rho'_i, \rho_j}}{1+2r_i^{\rho'_i, \rho_j}}s_i^S \right] + \beta(1 - m_j)s_i^P.
\end{aligned}$$

Player i 's best reply thus consists in choosing with probability one an action $a_i(s_i) = m_i s_i^P + (1 - m_i)s_i^S$ that is also linear in s_i with

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

After some algebra, one can verify that, for any cognitive profile $\rho = (\rho_i, \rho_j)$, the stage-2 game admits a unique linear continuation equilibrium and the latter is such that, for any $s_i = (s_i^P, s_i^S)$, σ_i^ρ , $i = 1, 2$, selects with certainty the action $a_i^\rho(s_i) = m_i^\rho s_i^P + (1 - m_i^\rho)s_i^S$, with

$$m_i^\rho = \frac{1 + 2r_j^\rho + r_i^\rho(1 - \beta) + 2r_i^\rho r_j^\rho(1 - \beta^2)}{1 + 2(r_i^\rho + r_j^\rho) + 4r_i^\rho r_j^\rho(1 - \beta^2)} \quad (16)$$

The above results also imply that, when the two players are expected to engage in cognition $\rho = (\rho_i, \rho_j)$ and, instead, player i selects cognition ρ'_i , in the stage-2 game, player i 's best response to player j following the strategy σ_j^ρ described above is to select, for any $s_i = (s_i^P, s_i^S)$, with certainty the action

$$a_i^{\rho'_i; \rho}(s_i) = m_i^{\rho'_i; \rho} s_i^P + (1 - m_i^{\rho'_i; \rho}) s_i^S$$

where

$$m_i^{\rho'_i; \rho} = \frac{1 + r_i^{\rho'_i, \rho_j}(1 + \beta) - 2\beta r_i^{\rho'_i, \rho_j} m_j^\rho}{1 + 2r_i^{\rho'_i, \rho_j}} \quad (17)$$

with m_j^ρ as in (16) but applied to player j .

Next, let $\eta_1 \equiv \phi_1 + \gamma_2$ and $\eta_2 \equiv \phi_2 + \gamma_1$ and 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 engage in cognition $\rho = (\rho_i, \rho_j)$ and, instead, player i selects cognition ρ'_i , player i 's ex-ante expected payoff (gross of the cognitive cost C_i but net of all terms that do not affect individual best responses) is equal to

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

Finally, 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

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

Using the characterization of the V_i functions above, after some algebra, we have that

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

Observe that UEC holds if $D \geq 0$ for $\rho' = (\hat{\rho}_i, \rho_j)$ and $\rho'' = (\rho_i, \rho_j)$, whereas ID holds if $D \geq 0$ for $\rho' = (\hat{\rho}_i, \hat{\rho}_j)$ and $\rho'' = (\hat{\rho}_i, \rho_j)$. In the proofs of Propositions 6 and 7 below we use this D function to establish all the results about UEC and ID in the two special cases of espionage and counter-espionage described in the main text.

Proof of Proposition 6. As explained in the main text, we abuse notation by letting each player's cognition coincides with the precision of the noise in the player's observation of the opponent's primary signal. Given any pair of cognitive levels $\hat{\rho}_i$ and ρ_i for player i , and any pair of cognitive profiles $\rho' = (\rho'_i, \rho'_j)$ and $\rho'' = (\rho''_i, \rho''_j)$, we then have that $r_i^{\rho_i, \rho'_j} = r_i^{\rho_i, \rho''_j} = \rho_i$, $r_i^{\hat{\rho}_i, \rho'_j} = r_i^{\hat{\rho}_i, \rho''_j} = \hat{\rho}_i$, $r_j^{\rho_i, \rho'_j} = r_j^{\hat{\rho}_i, \rho'_j} = \rho'_j$, $r_j^{\rho_i, \rho''_j} = r_j^{\hat{\rho}_i, \rho''_j} = \rho''_j$, $r_i^{\rho'} = \rho'_i$, $r_i^{\rho''} = \rho''_i$, $r_j^{\rho'} = \rho'_j$, and $r_j^{\rho''} = \rho''_j$.

Replacing the above formulas in the difference $D \equiv \left[V_i(\hat{\rho}_i; \rho') - V_i(\rho_i; \rho')\right] - \left[V_i(\hat{\rho}_i; \rho'') - V_i(\rho_i; \rho'')\right]$ in payoff differentials across the two profiles (see the characterization above, just before the beginning of the proof of the proposition), we then have that

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

Furthermore, in this case,

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

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,

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

To see that this game satisfies UEC, then 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 7. As explained in the main text, in this case, we abuse notation by letting each player's cognition coincide with the precision of the noise in the opponent's observation of the player's own primary signal. Given any pair of cognitive levels $\hat{\rho}_i$ and ρ_i for player i , and any pair of cognitive profiles $\rho' = (\rho'_i, \rho'_j)$ and $\rho'' = (\rho''_i, \rho''_j)$, we then have that $r_i^{\rho_i, \rho'_j} = r_i^{\hat{\rho}_i, \rho'_j} = \rho'_j$, $r_i^{\rho_i, \rho''_j} = r_i^{\hat{\rho}_i, \rho''_j} = \rho''_j$, $r_j^{\rho_i, \rho'_j} = r_j^{\rho_i, \rho''_j} = \rho_i$, $r_j^{\hat{\rho}_i, \rho'_j} = r_j^{\hat{\rho}_i, \rho''_j} = \hat{\rho}_i$, $r_i^{\rho'_i} = \rho'_i$, $r_i^{\rho''_i} = \rho''_i$, $r_j^{\rho'_j} = \rho'_j$, and $r_j^{\rho''_j} = \rho''_j$.

Replacing the above formulas into the function $D \equiv [V_i(\hat{\rho}_i; \rho') - V_i(\rho_i; \rho')] - [V_i(\hat{\rho}_i; \rho'') - V_i(\rho_i; \rho'')]$ introduced just before the beginning of the proof of Proposition 6, we have that, in this case,

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

Hence, we have that

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

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 $D \stackrel{\text{sgn}}{=} \beta(\hat{\rho}_i - \rho_i)H$, 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, $D \stackrel{\text{sgn}}{\equiv} \beta$. This game thus satisfies UEC for complements but not for substitutes.

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

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

Hence, $D \stackrel{\text{sgn}}{\equiv} -(\hat{\rho}_i - \rho_i)(\hat{\rho}_j - \rho_j)$. Independently of the sign of β , this game thus satisfies a form of negative ID. Q.E.D.