News and Archival Information in Games*

Ran Spiegler†

March 14, 2019

Abstract

I enrich the typology of players in the standard model of games with incomplete information, by allowing them to have incomplete "archival information" - namely, piecemeal knowledge of steady-state correlations among relevant variables. A player's type is defined by a conventional signal (a.k.a "news-information") as well as the novel "archive-information", formalized as a collection of subsets of variables. The player can only learn the marginal distributions over these subsets of variables. Drawing on prior literature on correlation neglect and coarse reasoning, I assume that the player extrapolates a well-specified probabilistic belief from his limited archival information according to the maximum-entropy criterion. This formalism expands our ability to capture strategic situations with "boundedly rational expectations."

---

*Financial support by ERC Advanced Investigator grant no. 692995 is gratefully acknowledged. I thank Kfir Eliaz and Heidi Thysen, as well as seminar audiences at Tel Aviv, Hebrew and New York Universities, for helpful comments. I also thank BRIQ for its hospitality during the writing of this paper.

†Tel Aviv University, University College London and CFM. URL: http://www.tau.ac.il/~rani. E-mail: rani@post.tau.ac.il.

1 Introduction

Conventional game theory distinguishes between two different kinds of information in static games. The first kind concerns the *realization* of exogenous variables - e.g. players' preferences. The second kind concerns the equilibrium statistical behavior of (exogenous as well as endogenous) variables - e.g., how players' bids in an auction vary with their preferences. Information of the first kind offers the player some knowledge regarding exogenous characteristics of the strategic situation; information of the second kind enables him to use this knowledge to draw inferences about other variables. To use a journalistic metaphor, the first kind is akin to a newsflash about a corporate scandal, whereas the second kind is what a reporter gets when he digs his newspaper's archives for evidence about the behavior of various actors in similar histotrial events. Accordingly, I refer to the two kinds of information as *news-information* and *archive-information*.

The standard model of Bayesian games offers a rich description of players' news-information, and leaves the task of describing players' archive-information to the solution concept. Introspective, "one-shot" solution concepts like rationalizability or level-k reasoning implicitly assume that players lack archive-information altogether. At the other extreme, Nash equilibrium presumes that players have complete archive-information - they know the steady-state joint distribution of all variables.

In the last two decades, the literature on equilibrium behavior with non-rational expectations proposed alternative solution concepts that retain the steady-state approach of Nash equilibrium, while replacing complete archive-information with a notion of limited feedback that players receive regarding the steady-state distribution - coupled with a model of how players form a belief given their feedback. E.g., in Jehiel (2005) players receive coarse feedback regarding the steady-state distribution, and extrapolate a belief that is measurable with respect to this feedback. I provide a detailed literature review in Section 5; at this stage, it suffices to say that virtually all prior proposals assume that players' feedback limitations are fixed. And neither provides a model of players' imperfect information (of either kind) regarding their opponents' archive-information.

Yet it is easy to think of real-life situations that call for such a model. For instance, sophisticated traders in financial markets may have accurate news-information *as well as* rich archive-information, whereas naive traders may be deprived of both. Thus, the two kinds of information may be correlated. We can also talk meaningfully about one player's news-information regarding another player's archive-information. To use a military-intelligence example, suppose that army 1 receives news from a dubious source that army 2 has just gained access to archival records of army 1's deployment of its forces as a function of, say, weather conditions. Or, consider the idea of one player having incomplete archive-information about another player's archive-information. For instance, army 1 may obtain a computer file that documents army 2's archival access in various types of situations. And one can imagine "high-order" elaborations of such statements, just as we do for news-information in the standard model of Bayesian games.

In this paper I present a new type space that combines both kinds of incomplete information in the context of static two-person games. A state of the world is described by the realization of a collection of variables. There is an objective prior distribution p over states. Following Aumann (1987), the description of a state also includes the realization of *endogenous* variables (players' actions, final allocations, etc.). Accordingly, I interpret p as a *steady-state distribution* in the system described by the model. The formalism also makes use of an explicit notational distinction between variables and the set of their labels L.

A player's type consists of two components. The first component - referred to as the player's N-information - is defined as the *realization* of a specific subcollection of the exogenous variables. This is the conventional notion of a signal: the player is informed of the realization of some exogenous variables. The novelty lies in the *second component* of the player's type - his R-information - which is defined as a subset of 2^L. It represents the player's "archival access" or "database privilege", and means that the player gets to learn the marginal of p over each of the variable collections defined by his R-information. Thus, rather than learning the entire joint distribution p (as in the case of Nash equilibrium), the player has piecemeal knowledge of it in the form of certain marginals. I assume that the player's payoff function is always measurable with respect to the variables about which he does get archival information.

The player forms a belief in two stages. First, he extrapolates a subjective probabilistic belief (defined over the variables about which he has data), thus forming a potentially distorted perception of the objective distribution p. The second stage is conventional: The player conditions his subjective belief (arrived at the end of the first stage) on his type via Bayes' rule, and thus forms a subjective belief over payoff-relevant outcomes as a function of his action. Equilibrium is defined in the familiar manner of trembling-hand perfection: If an action fails to maximize a player's subjective expected utility given his type, it is played with vanishing probability.

Of course, there are many extrapolation rules one could employ in the procedure's first stage. However, a recurring theme in the literature on equilibrium with non-rational expectations is that players apply some notion of *parsimony* when thinking about steady-state correlations - i.e., *they do not believe in correlations for which they lack direct evidence.* This tendency toward "correlation neglect" has been discussed extensively in the literature, both theoretically (e.g. Levy and Razin (2015)) and experimentally (e.g. Enke and Zimmermann (2017)). It has also been mentioned as a culprit in professional analysts' failure to predict major political and economic events.[1]

To capture this motive, I assume players use *maximum-entropy* extrapolation. That is, the player's belief is the distribution (over variables he has

[1] See Hellwig (2008) and https://fivethirtyeight.com/features/the-real-story-of-2016/.

archival data about) that maximizes entropy subject to being consistent with the marginals he learns. This subsumes existing notions in the literature as special cases, and is easy to calculate in certain applications.

1.1 An Example: The Prisoner's Dilemma

The following is a basic illustration of the formalism. Two players, denoted 1 and 2, play the following version of the Prisoner's Dilemma:

$a_1\backslash a_2$	C	D
C	$3,3$	$0,4$
D	$4,0$	$1,1$

There is no uncertainty regarding the game's payoff structure; the only uncertainty will be about players' archive-information.

Throughout the paper, I have in mind situations in which players lack an understanding of the game. Each player interacts *once*, after getting some feedback about the behavior of numerous generations of agents who assumed the players' current roles. If players knew they were playing a simultaneous-move game - let alone one in which players have a dominant action - they would use this knowledge to form beliefs. Instead, I assume that players' understanding of behavioral regularities is based *purely* on the learning feedback given by their archive-information. This is in the spirit of existing concepts like self-confirming, Berk-Nash or analogy-based expectations equilibrium (see Section 5).

Let R_i and a_i denote player i's R-information and action. As to players' news-information, each player learns the realization of his own R-information. Thus, player i's type is defined solely by R_i. A state of the world is thus described by the quadruple (R_1, R_2, a_1, a_2). The exogenous component of the prior p is the distribution over players' types $(p(R_1, R_2))$, whereas the endogenous components are the players' strategies given by the conditional distributions $(p(a_1 \mid R_1))$ and $(p(a_2 \mid R_2))$. Because players move simultaneously, p satisfies the conditional independence property $a_i \perp (R_j, a_j) \mid R_i$ for every $i = 1, 2$, $j \neq i$/(i.e., player i's action is independent of player j's type and action conditional on player i's own type).

The distribution over players' R-information is as follows. The set of variable labels is $L = \{l_{a_1}, l_{a_2}, l_{R_1}, l_{R_2}\}$. With probability $1 - \alpha$, both players have complete archive-information - that is, $R_1 = R_2 = \{L\}$. This means that players have full grasp of the steady-state distribution over all four variables. With probability α, players have incomplete archive-information: $R_1 = \{\{l_{R_1}, l_{a_1}\}, \{l_{a_1}, l_{a_2}\}\}$ and $R_2 = \{\{l_{R_2}, l_{a_2}\}, \{l_{a_1}, l_{a_2}\}\}$. That is, player i learns the joint steady-state distribution over his own archive-information and action, as well as the joint steady-state distribution over the action profile. The interpretation is as follows. The distribution over (a_1, a_2) represents a publicly available record of past realizations drawn from p, which enables players to learn the steady-state distribution over action profiles. In contrast, players' record of past realizations of the action and archive-information of agents who assumed their individual role is not always public.

Player i forms his belief in two stages. First, he extrapolates an unconditional subjective belief p_{R_i}, which is *defined over the variables he has data on*. Second, he conditions this belief on his type and action to evaluate the action - *just as he would in the standard model*! Thus, the only departure of this belief-formation procedure from the standard model lies in its first stage. Equilibrium is defined in the spirit of trembling-hand perfection (Selten (1975)), conventionally capturing the idea that long-run behavior involves infrequent blind experimentation. A profile of completely mixed strategies constitutes an ε-equilibrium if whenever $p(a_i \mid R_i) > \varepsilon$, a_i maximizes player i's expected utility with respect to his conditional subjective belief $p_{R_i}(a_j \mid R_i, a_i)$. An equilibrium is simply the limit of a sequence of ε-equilibria.

Let us derive players' beliefs as a function of their types. Complete archive-information induces rational expectations: $p_{\{L\}} = p$. Since p satisfies $a_j \perp a_i \mid R_i$, $p_{R_i}(a_j \mid a_i, R_i)$ is constant in a_i.[2] As a result, player i cannot think that his action affects a_j. And since C is strictly dominated, it follows that in any equilibrium, player i plays C with vanishing probability when $R_i = \{L\}$.

Now consider the realization $R_i = \{\{l_{R_i}, l_{a_i}\}, \{l_{a_1}, l_{a_2}\}\}$. The maximum-entropy extension of the marginals $p(R_i, a_i)$ and $p(a_1, a_2)$ is $p_{R_i}(R_i, a_1, a_2) = p(R_i, a_i)p(a_j \mid a_i)$. The belief satisfies $a_j \perp R_i \mid a_i$, capturing the notion of minimizing correlations beyond those the player learns.[3] Conditioning p_{R_i} on the player's type yields $p_{R_i}(a_j \mid R_i, a_i) = p(a_j \mid a_i)$. Thus, when $R_i = \{\{l_{R_i}, l_{a_i}\}, \{l_{a_1}, l_{a_2}\}\}$, player i forms a conditional belief *as if* he thinks that his action causes the opponent's action. He effectively mistakes the correlation between a_i and a_j (objectively a result of their respective dependence on players' correlated types) for a *causal effect* of the former on the latter. This should not be viewed as an explicit causal assumption that the player makes. Rather, it arises naturally from the (otherwise conventional) application of Bayesian conditioning to a belief that is parsimoniously extrapolated from observed correlations. Indeed, if the extrapolation were correct, there would be no fault with the effectively causal interpretation of $p(a_j \mid a_i)$.

We can now characterize the set of equilibria. One of them coincides with the standard game-theoretic prediction: Players always play D, regardless of their type. This equilibrium can be sustained by the following perturbation: Players choose C with small probability ε, independently of their type. Then, $p(a_j = D \mid a_i) = 1 - \varepsilon$ for all a_i. As a result, D is subjectively optimal for player i when $R_i = \{\{l_{R_i}, l_{a_i}\}, \{l_{a_1}, l_{a_2}\}\}$. We already saw that this is the case when $R_i = \{\{L\}\}$. Therefore, players' strategies constitute an ε-equilibrium.

Another equilibrium is for each player i to play D *if and only if* $R_i = \{L\}$.

[2] Because we examine completely mixed strategies, p has full support on (a_i, R_i), such that the conditional probability is well-defined.

[3] Note that p_{R_i} is not defined over R_j, because the player has no data about this variable.

We already saw that player i plays D when $R_i = \{L\}$; let us show that he plays C when $R_i = \{\{l_{R_i}, l_{a_i}\}, \{l_{a_1}, l_{a_2}\}\}$. Under the putative equilibrium, $p(a_j = C \mid a_i = C) = p(a_j = D \mid a_i = D) = 1$. Since player i's belief effectively interprets this perfect correlation between a_i and a_j causally - as if playing C (D) will cause player j to choose C (D) as well - C is a subjective best-reply.

There is a third, "hybrid" equilibrium, in which each player i plays C with probability $\lambda \in (0, 1)$ when $R_i = \{\{l_{R_i}, l_{a_i}\}, \{l_{a_1}, l_{a_2}\}\}$. Best-replying requires him to be indifferent between the two actions, given his conditional subjective belief:

$$3 \cdot p(a_j = C \mid a_i = C) = 4 \cdot p(a_j = C \mid a_i = D) + 1 \cdot p(a_j = D \mid a_i = D)$$

We can calculate

$$\begin{aligned} p(a_j &= C \mid a_i = C) = \frac{\alpha\lambda^2}{\alpha\lambda} = \lambda \\ p(a_j &= C \mid a_i = D) = \frac{\alpha\lambda(1-\lambda)}{1-\alpha\lambda} \end{aligned}$$

and obtain the solution $\lambda = 1/(3 - 2\alpha)$.

Thus, common variation in players' types gives rise to correlation between their actions; and in turn, one type's limited archive-information leads him to form an erroneous de-facto causal interpretation of this correlation. Of course, correlated actions could arise from other sources (e.g. correlated payoff shocks). What is special about the present example is that there are is no payoff uncertainty; the source of correlated actions is the very randomness of players' understanding of correlations. This "bootstrapping" effect will be a recurrent theme in this paper.

2 The Formalism

I restrict attention to simultaneous-move two-player interactions. The extension to more than two players is straightforward. Let X be a finite set of states of the world. This state space is endowed with a product structure, $X = \Theta \times S_1 \times S_2 \times \mathcal{R}_1 \times \mathcal{R}_2 \times A_1 \times A_2 \times Z$, where:

- $\theta \in \Theta$ is a state of Nature.
- $s_i \in S_i$ is player i's N-information (a conventional signal).
- $R_i \in \mathcal{R}_i$ is player i's R-information (to be endowed with explicit structure below).
- $a_i \in A_i$ is player i's action.
- $z \in Z$ is the game's outcome (e.g. an allocation of some resource).

Each one of these components may consist of a *collection* of variables (e.g., θ may have multiple dimensions, each described by a distinct variable). Some components may be suppressed in a given application (e.g., θ, s_1, s_2, z do not appear in the example of Section 1.1). I refer to the pair (s_i, R_i) as player i's *type*. Each player $i = 1, 2$ has a vNM utility function $u_i : X \to \mathbb{R}$. Let $p \in \Delta(X)$ be an objective prior distribution over the state space. Unless indicated otherwise, p has full support.

To reflect simultaneous moves, p satisfies the following conditional independence properties: $a_i \perp (\theta, s_j, R_j, a_j) \mid (s_i, R_i)$. That is, conditional on his type, player i's action is independent of the other exogenous variables and player j's action. It is also sensible (though not necessary) to assume that $z \perp (s_1, R_1, s_2, R_2) \mid (\theta, a_1, a_2)$. Thus, the exogenous components of p are the distribution over exogenous variables and the conditional distribution over outcomes (i.e., $(p(\theta, s_1, s_2, R_1, R_2))$ and $(p(z \mid \theta, a_1, a_2)))$; whereas the endogenous components are the players' strategies $\sigma_1 = (p(a_1 \mid s_1, R_1))$ and $\sigma_2 = (p(a_2 \mid s_2, R_2))$.

Comment: The notion of a state

A state of the world resolves *all* uncertainty, including the collection of endogenous variables a_1, a_2, z. Although unconventional, this approach has important precedents in the literature, notably in Aumann (1987), and it is fundamental to the present formalism. Accordingly, the prior p is interpreted as a *steady-state distribution* over *all* variables. I regard it as a representation of a long historical record of similar interactions. The individual game is a one-shot interaction between two players who move once against the historical background, and p records the collective experience of many other agents who assumed the players' roles in past interactions. An equally fundamental feature of the formalism is the definition of a state in terms of a collection of *variables*. Indeed, in Section 3.3 we will see an illustration that proliferation in the variables that define a state may have behavioral implications.

Finally, I impose explicit structure on R_i in order to substantiate its interpretation as archive-information. This will require a bit of notation. Enumerate the variables that define a state, such that $x = (x_1, ..., x_n)$, X_i is the set of values that variable x_i can take and i is the *label* of variable x_i. I will often use the alternative notation l_{x_i} for the label of x_i, in order to associate a variable with its label more transparently. Let $L = \{1, ..., n\}$ denote the set of *variable labels*. For every $B \subset L$, denote $x_B = (x_i)_{i \in B}$ and $X_B = \times_{i \in B} X_i$. Let $p^B \in \Delta(X_B)$ denote the marginal of p on X_B.

Armed with this notation, I assume that $\mathcal{R}_i \subset 2^L$. The meaning is that for any prior p, player i learns p^B for each $B \in R_i$. Thus, R_i represents player i's limited access to archival data - his "database privileges", so to speak. When $R_i = \{L\}$, player i has complete archive-information. In contrast, when R_i consists of a number of small subsets of L, player i's archive-information is incomplete. Let $N(R_i)$ denote the union of the members of R_i.

I impose two restrictions on $\mathcal{R}_i$. First, for every $R_i \in \mathcal{R}_i$ and every pair $B, B' \in R$, B and B' are not subsets of one another. Second, assume that u_i *is measurable with respect to* $x_{N(R_i)}$ for every $R_i \in \mathcal{R}_i$. That is, the player never lacks archival data about a variable that is necessary for defining his payoffs.

Belief formation

To make a decision, player i forms a probabilistic belief as a function of his type. In the standard model of Bayesian games, a player's type is defined solely in terms of his news-information and he forms his belief in a single step: Bayesian updating of p conditional on his type. In the present model, a player's type consists of two kinds of information, and so he forms his conditional probabilistic assessment in two stages, where each stage makes use of a different kind of information:

Stage one involves *maximum-entropy extrapolation from archive-information*: the player forms the unconditional belief $p_{R_i} \in \Delta(X_{N(R_i)})$ that solves

$$\max_{q \in \Delta(X_{N(R_i)})} \left[-\sum_{x_{N(R_i)}} q(x_{N(R_i)}) \ln(q(x_{N(R_i)})) \right] \tag{1}$$
$$\text{s.t. } q^B \equiv p^B \text{ for every } B \in R_i$$

That is, the player's unconditional belief over the variables about which he has archival data maximizes entropy subject to being consistent with the marginals his archival data enables him to learn. (The solution to the constrained maximization problem will always be unique.)

Stage two involves *conditioning on the player's type* and his action, according to conventional Bayesian updating. The player's conditional belief over $X_{N(R_i)}$ is thus $p_{R_i}(x_{N(R_i)} \mid s_i, R_i, a_i)$.[4]

[4] When player i lacks archival data about any of the variables s_i, R_i or a_i - i.e., $\{l_{a_i}, l_{s_i}, l_{R_i}\} \subsetneq N(R_i)$ - this variable can be safely omitted from the list of conditioned variables in.$p_{R_i}(x_{N(R_i)} \mid s_i, R_i, a_i')$.

Thus, each component of the player's type is associated with a particular *operation* that he performs on the objective prior p. The first stage involves *extrapolation*; the player's archive-information tells us what he extrapolates from. The second stage involves *conditioning*; the player's news-information tells us what he conditions on. This stage utilizes the canonical rule of Bayesian updating. By comparison, there is no "canonical" extrapolation rule. Nevertheless, there is a common intuition that extrapolating a belief from partial data should follow some *parsimony* criterion. A number of solution concepts in the literature (see Section 5) involve parsimonious treatment of *correlations*. This is a plausible criterion for extrapolating beliefs from partial learning feedback when players lack sufficient understanding of the game (and when variables lack structure that would suggest other extrapolation criteria, such as monotonicity). The maximum-entropy criterion (which originates from statistical physics and has a rich tradition in data analysis (see Jaynes (1957)) systematizes this idea. It regards minimal assumptions on correlations as parsimonious, and thus looks for the distribution that exhibits maximal statistical independence subject to being consistent with observed correlations.

Equilibrium

Having defined players' beliefs as a function of their types, we are ready to introduce the notion of equilibrium, which is a standard trembling-hand perfection concept.

Definition 1 *Fix $\varepsilon > 0$ and the exogenous components of the prior p. A profile of full-support strategies (σ_1, σ_2) is an ε-equilibrium if for every $i = 1, 2$ and every a_i, s_i, R_i for which $p(a_i \mid s_i, R_i) > \varepsilon$,*

$$a_i \in \arg\max_{a_i'} \sum_{x_{N(R_i)}} p_{R_i}(x_{N(R_i)} \mid s_i, R_i, a_i') u_i(x_{N(R_i)})$$

A strategy profile (σ_1^*, σ_2^*) (which need not satisfy full support) is an equilib-

rium if it is the limit of a sequence of ε-equilibria with $\varepsilon \to 0$. We say in this case that p is in equilibrium.

Establishing existence of equilibrium is straightforward. Because p_{R_i} is a continuous function of p, the proof is essentially the same as in the case of standard trembling-hand perfect equilibrium.

Closed forms of p_R

In some cases of interest, the maximum-entropy extension of a given collection of marginal distributions takes a tractable, interpretable form. Indeed, some of these cases are familiar from the literature on equilibrium models with non-rational expectations. The simplest case is where R_i consists of a single subset $\{B\}$, $B \subset L$. In this case, $p_{R_i} = p^B$. This is trivially pinned down by the requirement that the player's belief is consistent with p^B. Though apparently trivial, this case can have significant behavioral implications. Piccione and Rubinstein (2003), Eyster and Piccione (2013) and Eliaz et al. (2018) study interactive models in which agents' beliefs can be described in these terms. In macroeconomics, such beliefs appear in so-called "restricted perceptions equilibrium" (see Woodford (2013)).

For a slightly more involved example that is structurally the same as in Section 1.1, let $x = (\theta^1, \theta^2, R_1, R_2, a_1, a_2)$, where (θ^1, θ^2) is the state of Nature. The realization $R = \{\{l_{\theta^1}, l_{\theta^2}\}, \{l_{\theta^2}, l_{a_2}\}\}$ indicates that player 1 learns the distribution of the state of Nature as well as the joint distribution of player 2's action and a coarse description of the state of Nature (given by the component θ^2). The maximum-entropy extension of these marginals is $p_R(\theta^1, \theta^2, a_2) = p(\theta^1, \theta^2)p(a_2 \mid \theta^2)$. This is what the notion of *analogy-based expectations* in static games (Jehiel and Koessler (2008)) would prescribe when θ^2 is defined as the analogy class to which (θ^1, θ^2) belongs. When we omit θ^2 from the model such that $R = \{\{l_{\theta^1}\}, \{l_{a_2}\}\}$, we obtain $p_R(\theta^1, a_2) = p(\theta^1)p(a_2)$ - i.e. a belief that a_2 is independent of the state of Nature. This is an instance of "fully cursed" beliefs (Eyster and Rabin (2005)).

Mailath and Samuelson (2018) analyze information aggregation when agents follow a similar belief-formation model. The relevant variables for an individual player are the consequence variable z, a collection of variables x_M that he regards as sufficient predictors of z, and a profile b of other players' beliefs. The player's subjective belief over these variables is defined as $p(b)p(x_M \mid b)p(z \mid x_M)$, which is a way of writing the maximum-entropy extension of the marginals $p(b, x_M)$ and $p(x_M, z)$.

Thus, maximum-entropy extrapolation can be regarded as a principle that unifies a number of precedents in the literature. These examples and the others I examine in this paper fall into the following category.

Definition 2 *We say that R satisfies the running intersection property (**RIP**) if its elements can be ordered $B^1, ..., B^m$ such that for every $k = 2, ..., m$, $B^k \cap (\cup_{j<k} B^j) \subseteq B^i$ for some $i = 1, ..., k-1$.*

It holds trivially for $m = 2$. The $m = 3$ collection $\{\{1,2\},\{2,3\},\{3,4\}\}$ satisfies the property, whereas the $m = 3$ collection $\{1,2\},\{2,3\},\{1,3,4\}\}$ violates it. RIP ensures a simple closed form for the maximum-entropy extension of the marginals $(p^B)_{B\in R}$.

Proposition 1 (Hajek et al. (1992)) *When R satisfies RIP, the maximum-entropy extension of $(p^B)_{B\in R}$ is given by*

$$p_R(x_{N(R)}) = \prod\nolimits_{B^1,...,B^m} p(x_{B^k - (\cup_{j<k} B^j)} \mid x_{B^k \cap (\cup_{j<k} B^j)}) \qquad (2)$$

where the enumeration $1, ..., m$ validates RIP.

For instance, when $R = \{\{1,2\},\{2,3\},\{3,4\}\}$,

$$\begin{aligned} p_R(x_1, x_2, x_3, x_4) &= p(x_1, x_2)p(x_3 \mid x_2)p(x_4 \mid x_3) \qquad (3) \\ &= p(x_1)p(x_2 \mid x_1)p(x_3 \mid x_2)p(x_4 \mid x_3) \end{aligned}$$

Thus, RIP allows us to write p_R as a factorization of $p(x_{N(R)})$ into marginal and conditional distributions. Moreover, the factorization has a *causal* interpretation. For instance, (3) looks as if p_R is consistent with the causal chain $1 \to 2 \to 3 \to 4$. This is a general property of (2). This formula can be rewritten as a factorization of p according to a *directed acyclic graph* whose set of nodes is $N(R)$, such that R is the set of maximal cliques in the graph's non-directed version. Thus, when R satisfies RIP, it has a dual interpretation as the player's subjective causal model, such that p_R is the outcome of fitting his model to long-run data. Indeed, RIP is a familiar concept in the literature on graphical probabilistic models (see Cowell et al. (1999)). I do not elaborate on the causal interpretation in this paper and refer the reader to Spiegler (2016,2017,2018) for more details.

Can players' beliefs be "rationalized"?

In some cases, the beliefs generated by partial R-information can be replicated by a conventional model with complete archive-information and limited news-information. For example, let $x = (a_1, s_1, \theta, s_2, a_2)$. Suppose that the signals s_1 and s_2 are independent conditional on θ (and by assumption, players' actions are independent of θ conditional on their signal). Players' R-information is constant and therefore can be omitted from the description of the state of the world. Suppose that player 1's R-information is $\{\{\delta_{a_1}, \delta_{s_1}\}, \{\delta_\theta, \delta_{a_2}\}\}$. Then, $p_{R_1}(\theta, a_2 \mid s_1, a_1) = p(\theta, a_2)$. The same subjective belief would arise in a standard rational-expectations model in which player 1 is entirely *uninformed* of θ.

However, this is an atypical example. In the same setting, let player 1's R-information be $\{\{\delta_{a_1}, \delta_{s_1}, \delta_\theta\}, \{\delta_{a_2}\}\}$. Then, $p_{R_1}(\theta, a_2 \mid s_1, a_1) = p(\theta \mid s_1)p(a_2)$. That is, the player forms rational expectations regarding θ but fails to perceive the dependence of a_2 and θ. This effect cannot be reproduced by a standard model: If a_2 and θ are correlated and player 1 can draw inferences about θ from his signal s_1, then in a standard model he will also draw inferences about a_2. Observations in a similar vein were made by Eyster and Rabin (2005) and Jehiel and Koessler (2008).

The following example exhibits a different kind of inconsistency with a standard model. Suppose that player 2 has a degenerate action set, such that the game is reduced to a single-agent decision problem. Let $x = (a_1, s_1^1, s_1^2, \theta)$, where s_1^1 and s_1^2 are two components of player 1's N-information. Suppose that his R-information is $\{\{a_1\}, \{\theta, s_1^1\}, \{\theta, s_1^2\}\}$. Then, $p_{R_1}(\theta, s_1^1, s_1^2) = p(\theta)p(s_1^1 \mid \theta)p(s_1^2 \mid \theta)$. It is easy to construct objective distributions $(p(\theta, s_1^1, s_1^2))$ such that the distribution over player 1's subjective conditional belief $p_{R_1}(\theta \mid s_1^1, s_1^2)$ will violate Bayes-plausibility - i.e., $\sum_{s_1^1, s_1^2} p(s_1^1, s_1^2) p_{R_1}(\theta \mid s_1^1, s_1^2)$ will not coincide with $p(\theta)$. Such a distribution over subjective beliefs can never be sustained by a standard rational-expectations model.[5]

Maximum entropy and the insufficient-reason principle

To see the first stage of players' belief-formation procedure from a slightly different angle, let q^* denote the uniform distribution over $X_{N(R_i)}$. Consider a reformulation of the procedure's first stage, which changes the objective function in (1) into

$$\sum_{x_{N(R_i)}} q(x_{N(R_i)}) \frac{\ln(q(x_{N(R_i)}))}{\ln(q^*(x_{N(R_i)}))} \tag{4}$$

This formulation is equivalent because q^* is uniform and therefore $\ln(q^*(z))$ is a constant. Expression (4) is the *relative-entropy distance* (a.k.a Kullback-Leibler Divergence) of q from q^*. The extrapolated belief p_{R_i} minimizes this distance from the uniform distribution, out of all the distributions in $\Delta(X_{N(R_i)})$ that are consistent with $(p^B)_{B \in R_i}$. The interpretation is that the player initially has a uniform ("Laplacian") prior over X, and then he learns the true marginals $(p^B)_{B \in R_i}$. When these marginals refute the Laplacian belief, the player revises his initial theory in a minimalistic fashion that is captured by relative-entropy distance. Indeed, in the absence of any data, maximum entropy would imply a Laplacian belief. Thus, maximum entropy can be viewed as an extension of the principle of insufficient reason.

[5] See a related discussion in Section 5 of Spiegler (2018).

3 Information about Archive-Information

This section presents three illustrations of the formalism's ability to describe novel yet realistic notions of "information about an opponent's information". For expositional clarity, I fix the conventional aspects and vary the novel ones. The game that serves as my template is familiar from the "global games" literature since Rubinstein (1989), Carlsson and van Demme (1993) and Morris and Shin (1998). Its structure makes high-order strategic reasoning crucial for players' behavior, and therefore enables us to illustrate the novel kinds of high-order reasoning that the formalism can capture. In particular, we will see the failures to coordinate that arise from players' limited news-information or archive-information regarding the opponent's archive-information.

Thus, throughout the section, I examine a 2×2 game in which a_1 and a_2 take values in $\{0, 1\}$ and the payoff matrix is

$a_1 \backslash a_2$	1	0
1	$\delta\theta, \delta\theta$	$-1, 0$
0	$0, -1$	$0, 0$

where $\delta \in (0, 1)$ is a constant, $\theta \in \{0, 1\}$ is the state of Nature, and $p(\theta = 1) = \frac{1}{2}$. When θ is common knowledge and players have rational expectations, they both find $a = 0$ a strictly dominant action when $\theta = 0$, whereas under $\theta = 1$ they know they are playing a coordination game with two Nash equilibria: $(0, 0)$ and $(1, 1)$.

3.1 News-Information about Archive-Information

Suppose that player 1 has complete R-information with probability one. Therefore, we can omit R_1 as a variable from the description of the state of the world. Player 2's R-information is distributed independently of θ. With probability $\alpha \in (0, \frac{1}{2})$, $R_2 = \{\{l_\theta, l_{a_1}, l_{a_2}\}\}$; I give this realization a shorthand notation $R_2 = 1$. With the remaining probability $1-\alpha$, $R_2 = \{\{l_\theta\}, \{l_{a_1}\}, \{l_{a_2}\}\}$ - a realization also denoted $R_2 = 0$. Thus, R_2 records whether player 2 learns the correlation between the state of Nature and players' actions.

As to players' N-information, both players perfectly learn the realization of θ. In addition, player 1 receives a signal $s_1^{arch} \in \{0, 1\}$ regarding the value of R_2. Assume $p(s_1^{arch} = R_2) = q$ for every R_2, independently of θ, where $q \in (\frac{1}{2}, 1)$. Player 1's N-information is thus (θ, s_1^{arch}). Player 2's N-information consists of θ alone, such that there is no need to include a distinct variable s_2 in the definition of the state of the world: $x = (\theta, s_1^{arch}, R_2, a_1, a_2)$.

Let us construct player 2's first-stage (unconditional) belief as a function of his R-information: $p_{R_2=1}(\theta, a_1, a_2) = p(\theta, a_1, a_2)$ and $p_{R_2=0}(\theta, a_1, a_2) = p(\theta)p(a_1)p(a_2)$. The derivations correspond to special cases presented in Section 2. It follows that player 2's belief conditional on his type and action is

$$
\begin{aligned}
p_{R_2=1}(a_1 \mid \theta, R_2, a_2) &= p(a_1 \mid \theta) \\
p_{R_2=0}(a_1 \mid \theta, R_2, a_2) &= p(a_1)
\end{aligned}
$$

The realization $R_2 = 0$ captures a "fully cursed" player (as in Eyster and Rabin (2005)) who does not perceive the correlation between player 1's action and the state of Nature. Following Ettinger and Jehiel (2010), this case can also be interpreted as a situation in which player 2 commits the *Fundamental Attribution Fallacy* - i.e., he does not realize that player 1's behavior can be influenced by the state of Nature. More concretely, imagine that players' dilemma is whether to be nice in a social situation. When $\theta = 0$, players cannot afford to be nice. When $\theta = 1$, there are gains from mutually nice behavior. Player 1's action is potentially responsive to the social situation. However, when $R_2 = 0$, player 2 lacks access to the record of player 1's past behavior and does not get to learn this correlation; he extrapolates a belief that treats player 1's behavior as a non-situational statistical pattern.

While the realization $R_2 = 1$ does not exhibit the fundamental attribution error, it does not induce rational expectations. Rather, it captures a *"second-order" attribution error*: the player fails to realize that player 1's behavior is responsive to his news-information s_1^{arch} about whether player 2 exhibits the fundamental ("first-order") attribution error.

Proposition 2 *There is a unique equilibrium in this example: Both players always play* $a = 0$.

Proof. First, when $\theta = 0$, both players choose $a = 0$. To see why, note first that player 1 as well as player 2 under $R_2 = 1$ have rational expectations, and therefore correctly recognize that $a = 0$ is a dominant action under $\theta = 0$. Second, when $R_2 = 0$, we saw that player 2 believes that player 1 mixes over actions independently of θ, R_2, a_2. Since $p(\theta = 1) = \frac{1}{2}$, the previous paragraph implies that $p(a_1 = 1) \leq \frac{1}{2}$. Therefore, player 2's expected utility from $a_2 = 1$ against his subjective belief is $a_2 = 0$, regardless of the value of θ. We have thus established that $a_2 = 0$ in equilibrium whenever $\theta R_2 = 0$.

Let us try to sustain an equilibrium in which $p(a_1 = 1 \mid \theta = 1) > 0$. First, derive player 1's posterior belief regarding R_2 as a function of his signal s_1^{arch}:

$$p(R_2 = 1 \mid s_1^{arch} = 1) = \frac{\alpha q}{\alpha q + (1-\alpha)(1-q)}$$

$$p(R_2 = 1 \mid s_1^{arch} = 0) = \frac{\alpha(1-q)}{\alpha(1-q) + (1-\alpha)q}$$

Recall that player 2 plays $a_2 = 0$ whenever $\theta R_2 = 0$. By our assumptions on α and q, $p(R_2 = 1 \mid s_1 = 0) < \frac{1}{2}$. Therefore, when player 1 observes $s_1^{arch} = 0$, his unique best-reply is $a_1 = 0$. It follows that

$$p(a_1 = 1 \mid \theta = 1) \leq p(s_1 = 1) = \alpha q + (1-\alpha)(1-q) < \frac{1}{2}$$

This means that player 2's best-reply is $a_2 = 0$, regardless of R_2. Player 1's best-reply is necessarily $a_1 = 0$ regardless of s_1, a contradiction. It follows that player 1 always plays $a_1 = 0$ in any equilibrium. Completing the proof is now straightforward. ■

Player 2's "second-order attribution error" is the key to this negative result. If $R_2 = 1$ represented complete R-information, player 2 would be able to infer from his own archive-information (if q is high enough) that player 1 is likely to observe $s_1^{arch} = 1$ and play $a_1 = 1$, such that player 2's best-reply would be $a_2 = 1$. In contrast, our definition of $R_2 = 1$ means that player 2 effectively fails to condition his forecast of a_1 on R_2. As a result, he ends up underestimating the conditional probability that player 1 will choose $a_1 = 1$.

3.2 Archive-Information about Archive-Information

In this sub-section, I analyze a situation in which equilibrium patterns arise from players' random archive-information about their opponent's archive-information. We already saw a simple example of this phenomenon in Section 1.1, where players' archive-information sometimes failed to include data about this variable. The following example is a more elaborate variant on this theme.

Unlike the previous sub-section, this example treats players symmetrically: R_1 and R_2 both get two values, denoted 0 and 1. Assume that R_1 and R_2 are distributed *uniformly* and independently of θ. Moreover, they are perfectly correlated: $R_1 = 1$ if and only if $R_2 = 1$. Players' news-information consists of θ and R. There is no need to specify a distinct news-information variable, such that the state of the world x can be written as $x = (\theta, R_1, R_2, a_1, a_2)$. Players' R-information is given explicitly as follows:

$$\begin{aligned} R_i &= 0 : \{\{l_\theta, l_{a_1}\}, \{l_\theta, l_{a_2}\}\} \\ R_i &= 1 : \{\{l_\theta, l_{a_1}\}, \{l_\theta, l_{a_2}\}, \{l_{R_i}, l_{R_j}\}, \{l_{R_j}, l_{a_j}\}\} \end{aligned}$$

Thus, player i's type is defined by whether he has archive-information about player j's archive-information. Specifically, $R_i = 0$ means that player

i only learns the joint distribution of individual players' actions with the state of Nature, whereas $R_i = 1$ means that the player also learns how player j's archive-information correlates with his action as well as with player i's own archive-information. Both realizations of R_i satisfy RIP, such that

$$\begin{aligned} p_{R_i=0}(\theta, a_i, a_j) &= p(\theta, a_i)p(a_j \mid \theta) \\ p_{R_i=1}(\theta, a_i, a_j, R_i, R_j) &= p(\theta, a_i)p(a_j \mid \theta)p(R_j \mid a_j)p(R_i \mid R_j) \end{aligned}$$

In both cases, p_{R_i} treats a_j as independent of a_i conditional on θ, R_i. Therefore, we can omit a_i and work with the simpler forms:

$$\begin{aligned} p_{R_i=0}(\theta, a_j) &= p(\theta)p(a_j \mid \theta) \\ p_{R_i=1}(\theta, a_j, R_i, R_j) &= p(\theta)p(a_j \mid \theta)p(R_j \mid a_j)p(R_i \mid R_j) \end{aligned}$$

In this example, there are equilibria in which players' behavior is independent of their R-information. The reason is that if a_j is independent of R_j, the realizations $R_i = 1$ and $R_i = 0$ induce the same $p_{R_i}(a_j \mid \theta, R_i, a_i)$, and therefore there is no reason for player i to vary his action with R_i. Let us examine whether there are equilibria that allow players' actions to vary with their archive-information.

Proposition 3 *There is a symmetric equilibrium in which* $p(a_i = 1 \mid \theta, R_i) = \theta R_i$ *for each* i*, if and only if* $\delta \geq \frac{1}{3}$*. There exist no additional symmetric pure-strategy equilibria in which* a_i *varies with* R_i*.*

Proof. In any equilibrium, players choose $a = 0$ whenever $\theta = 0$, independently of R_i. The reasoning is the same as in the previous sub-section, and therefore omitted here. Let us now derive player i's conditional belief over a_j conditional on $\theta = 1$ and each of the two realizations of R_i:

$$p_{R_i=0}(a_j = 1 \mid \theta = 1, R_i = 0) = p(a_j = 1 \mid \theta = 1)$$

and

$$\begin{aligned}
p_{R_i=1}(a_j &= 1 \mid \theta = R_i = 1) = \frac{p_{R_i=1}(\theta = R_i = a_j = 1)}{\sum_{a_j'} p_{R_i=1}(\theta = R_i = 1, a_j')} \\
&= \frac{p(\theta = 1) p(a_j = 1 \mid \theta = 1) \sum_{R_j} p(R_j \mid a_j') p(R_i = 1 \mid R_j)}{p(\theta = 1) \sum_{a_j'} p(a_j' \mid \theta = 1) \sum_{R_j} p(R_j \mid a_j') p(R_i = 1 \mid R_j)} \\
&= \frac{p(a_j = 1 \mid \theta = 1) p(R_j = 1 \mid a_j = 1)}{p(a_j = 1 \mid \theta = 1) p(R_j = 1 \mid a_j = 1) + p(a_j = 0 \mid \theta = 1) p(R_j = 1 \mid a_j = 0)}
\end{aligned}$$

Suppose that players condition their (symmetric, pure-strategy) equilibrium action on their R-information. Then, $p(a_j = 1 \mid \theta = 1) = \frac{1}{2}$. Therefore, when $R_i = 0$, player i's best-reply is $a_i = 0$. It follows that if we wish to sustain a symmetric pure-strategy equilibrium in which a_i varies with R_i, it must be the case that $p(a_j = 1 \mid \theta = 1, R_j) = R_j$. We can now calculate $p_{R_i=1}(a_j = 1 \mid \theta = R_i = 1)$, by plugging the terms

$$p(a_j = 1 \mid \theta = 1) = \frac{1}{2} \qquad p(R_j = 1 \mid a_j = 1) = 1$$

and

$$\begin{aligned}
p(R_j &= 1 \mid a_j = 0) = \frac{p(R_j = 1, a_j = 0)}{p(a_j = 0)} \\
&= \frac{p(R_j = 1) p(\theta = 0)}{p(\theta = 0) + p(\theta = 1) p(R_j = 0)} = \frac{\frac{1}{2} \cdot \frac{1}{2}}{\frac{1}{2} + \frac{1}{2} \cdot \frac{1}{2}} = \frac{1}{3}
\end{aligned}$$

Thus,

$$p_{R_i=1}(a_j = 1 \mid \theta = R_i = 1) = \frac{\frac{1}{2} \cdot 1}{\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \frac{1}{3}} = \frac{3}{4}$$

Therefore, in order for $a_i = 1$ to be a best-reply when $\theta = R_i = 1$, it must be the case that

$$\frac{3}{4} \cdot \delta - \frac{1}{4} \cdot 1 \geq 0$$

hence, $\delta \geq \frac{1}{3}$ is a necessary and sufficient condition for an equilibrium in which $a_i = \theta R_i$ for every player i. Any other symmetric pure-strategy equilibrium exhibits R-independent actions. ■

Thus, as long as the gains from good coordination in $\theta = 1$ are sufficiently high, it is possible to sustain an equilibrium in which players coordinate efficiently if and only if they have rich archive-information. No other pattern of correlation between players' actions and their archive-information is sustainable in symmetric pure-strategy equilibrium.

Importantly, the requirement $\delta \geq \frac{1}{3}$ in this result arises from the fact that $R_i = 1$ represents *partial* archive-information. In this case, player i learns the pairwise correlation of a_j with θ and R_j; his failure to learn the *joint* correlation of a_j with θ, R_j limits the extent to which he updates his belief over a_j. By comparison, consider the alternative specification in which $R_i = 1$ represents complete archive-information. Then, the above equilibrium can be sustained for *any* $\delta > 0$.

3.3 Hierarchical Archive-Information

The representation of a state of the world in terms of a collection of variables is fundamental to the formalism. Furthermore, unlike standard Bayesian games, collapsing the collections θ, s_i or z into a single variable is not innocuous. For instance, let $\theta = (\theta^1, ..., \theta^K)$ and assume that there exists no $B \in R_i$ that contains $\{l_{\theta^1}, ..., l_{\theta^K}\}$. The belief distortions that arise from player i's partial R-information in this case cannot be reproduced if we collapse θ into a single variable.

Perhaps the most interesting case of this effect is where *players' archive-information itself* is represented by a collection of variables. In the examples of Sections 1.1 and 3.2, a player's R-information was one of the variables about which the other player had R-information. The formalism's capacity for such cross-references is one of its prime virtues - in analogy to the Bayesian-game formalism's ability to describe one player's N-information regarding another player's N-information. And as in the case of Bayesian games, it is natural to think of *hierarchical constructions* of this inter-dependence.

The starting point of a hierarchical definition of players' R-information is a collection of *basic* variables. Let $B \subset L$ be the set of labels of the basic variables. These would include variables that define the state of Nature, players' news-information and actions, as well as consequence variables. For each player i, there is a collection of variables $R_i^1, ..., R_i^m$, $m \geq 2$, where $R_i^1 \subset 2^B$; $R_i^2 \subset 2^{B \cup \{R_i^1, R_j^1\}}$; and for every $k = 3, ..., m$,

$$R_i^k \subset 2^{B \cup \{l_{R_i^h}, l_{R_j^h}\}_{h=1,...,k-1}}$$

In addition, every element in R_i^k, $k \geq 2$, includes $l_{R_i^{k-1}}$ or $l_{R_j^{k-1}}$. Define $R_i = \cup_{k=1,...,m} R_i^k$.

The interpretation of this hierarchical construction is as follows: R_i^1 is the player's first-order archive-information, describing his knowledge of correlations among basic variables; R_i^2 is the player's second-order archive-information, describing his knowledge of how players' first-order archive-information is correlated with the basic variables; and so forth.

The following is a simple example of hierarchically defined R-information in the context of the coordination game of this section. Players have common R-information, which is distributed independently of θ. Players' N-information coincides with θ - i.e., they perfectly learn the state of Nature. The basic variables are θ, a_1, a_2. Suppose that for every $k = 1, 2, ...$, R^k takes two values, given the shorthand notation 0 and 1 and defined explicitly as follows:

k	$R^k = 0$	$R^k = 1$
1	$\{\{l_\theta\}, \{l_{a_1}\}, \{l_{a_2}\}\}$	$\{\{l_\theta, l_{a_1}, l_{a_2}\}\}$
2	$\emptyset$	$\{\{l_\theta, l_{a_1}, l_{a_2}, l_{R^1}\}\}$
3	$\emptyset$	$\{\{l_\theta, l_{a_1}, l_{a_2}, l_{R^1}, l_{R^2}\}\}$
$\vdots$	$\vdots$	$\vdots$
m	$\emptyset$	$\{\{l_\theta, l_{a_1}, l_{a_2}, l_{R^1}, ..., l_{R^{m-1}}\}\}$
$\vdots$	$\vdots$	$\vdots$

The only values of R that are realized with positive probability are those for which $R^k = 1$ implies $R^{k-1} = 1$, for every $k > 1$. Therefore, it is convenient to represent R by the largest number k for which $R^k = 1$. Specifically, let $p(R = k) = \gamma(1-\gamma)^k$ for every $k = 0, 1, ...$ Note that $R = k$ means that players perceive actions as a function of $\theta, R^1, ..., R^{k-1}$.

Proposition 4 *Suppose that $\gamma > \frac{1}{2}$. Then, there is a unique equilibrium, in which players always play $a = 0$.*

Proof. The proof is by induction on k. As a first step, observe that by the same argument as in previous sub-sections, $a_i = 0$ whenever $\theta R^1 = 0$. Suppose that we have shown that $a_i = 0$ when $\theta = 1$ and $R < k$, and consider the case of player 1, say, when $\theta = 1$ and $R = k$. The player will find it optimal to play $a_1 = 1$ only if $p_{R=k}(a_2 = 1 \mid \theta = 1, R = k, a_1 = 1) > \frac{1}{2}$. However,

$$\begin{aligned} p_{R=k}(a_2 &= 1 \mid \theta = 1, R = k, a_1) = p(a_2 = 1 \mid \theta = R^1 = \cdots = R^{k-1} = 1) \\ &\leq \frac{p(R \geq k)}{p(R \geq k-1)} = 1 - \gamma < \frac{1}{2} \end{aligned}$$

and therefore, the player's best-reply is $a_1 = 0$. ■

The intuition for this result is simple. When players have $R = k$, they only perceive correlations between actions and R-information of order $k - 1$

and below. By the assumption that $\gamma > \frac{1}{2}$, players are more likely to lack R-information of order k conditional on having R-information of order $k-1$, and by the inductive step, they play $a = 0$ in that case. It follows that when a player has $R = k$, he believes it is more likely that the opponent will play $a = 0$, hence the best-reply is to play $a = 0$, too.

4 An Extended Example: "Market Savvy"

In this section I present a more elaborate example that further demonstrates the formalism's scope and illustrates a few of its novel features. The example addresses a fundamental economic question: What determines the performance of market agents? In particular, do "savvier" traders earn higher profits?[6] The example highlights two aspects of "market savvy", which correspond to news- and archive-information: accurate information about trade opportunities, and correct assessment of the extent to which competing traders chase (and therefore dissipate) such opportunities when they arise. The correlation between these two aspects of market savvy turns out to be a key factor.

Let (ψ, θ) be a state of Nature that is drawn uniformly from $[0,1] \times \{0,1\}$, where $\psi \in [0,1]$ is the "*location*" in which a trade opportunity can arise, and $\theta \in \{0,1\}$ indicates whether such an opportunity actually exists. Traders' R-information is distributed independently of the state of Nature (I describe the exact distribution below). Given (ψ, θ, R_i), trader i receives a signal $s_i = \psi$ with probability $\theta q_{R_i} + (1-\theta)(1-q_{R_i})$ and $s_i = \emptyset$ with the remaining probability, where $q_{R_i} \in (\frac{1}{2}, 1)$ is a constant that characterizes the type R_i. The larger q_{R_i}, the higher the probability that the trader learns the location of a market opportunity when this knowledge is relevant. Having received

[6]Eyster and Piccione (2013) explored this question in the context of a competitive asset market model. Applying this paper's terminology retroactively, traders in their model are characterized by incomplete archive-information, given by a single variable set that differs across trader types.

the signal s_i, trader i chooses a location $a_i \in [0,1]$. Define two induced binary variables. First, $t_i = 1$ (0) when $s_i \neq \emptyset$ ($s_i = \emptyset$). That is, t_i indicates whether trader i is informed of ψ. Second, $z_i = \mathbf{1}(a_i = \psi)$. That is, z_i indicates whether the trader chooses the precise location in which a market opportunity could arise. The trader's payoff is

$$\theta z_i(3 - 2z_i - 2z_j)$$

For a concrete story behind this model, think of locations as financial securities whose fundamental value is normalized to zero. Virtually all securities are fairly priced and provide no arbitrage opportunities. Only one security ψ is mispriced and has a "thin" market, such that traders who take a position in this security affect its price. The term $3-2z_i-2z_j$ represents the security's market-clearing price. If only one trader takes a position in this security he earns a net gain, but if both traders take a position they make a net loss. An individual trader's sophistication thus consists of identifying the needle-in-a-haystack arbitrage opportunity (news-information) and assessing the likelihood that the other trader fails to spot an arbitrage opportunity when it arises (archive-information).

The steady-state distribution p is defined over the state of the world $x = (\theta, \psi, R_1, R_2, s_1, s_2, t_1, t_2, a_1, a_2, z_1, z_2)$. I focus on *symmetric* equilibria - i.e., each trader i plays the same strategy $\sigma_i = (p(a_i \mid s_i, R_i))$. When $\theta = 0$, trader i's payoff is zero, regardless of his action. When $s_i = \emptyset$, trader i is uninformed of ψ. Since ψ is drawn from an atomless distribution, any action $a_i \in [0,1]$ induces $z_i = 0$ (and therefore zero payoff) with probability one when $s_i = \emptyset$. It follows that trader i's action only matters when $\theta = t_i = 1$. Choosing $a_i \neq \psi$ in this case induces $z_i = 0$ (and therefore zero payoff) with probability one. If, on the other hand, the trader chooses $a_i = \psi$ (such that $z_i = 1$), his payoff is $\theta(1 - 2z_j)$.

It follows that trader i's best-replying behavior is reduced to the following simple problem. When $t_i = 0$, trader i is constrained to satisfy $z_i = 0$ with probability one. And when $t_i = 1$, he effectively chooses $z_i \in \{0,1\}$ to maximize

$$z_i[3 - 2z_i - 2p_{R_i}(z_j = 1 \mid \theta = 1)]$$

This implies the following preliminary characterization of symmetric equilibria in this model.

Claim 1 *In any symmetric equilibrium,* $p(z_i = 1 \mid t_i = 0, R_i) = 0$ *for any* R_i*; and*

$$\begin{array}{lcl} p_{R_i}(z_j = 1 \mid \theta = 1) < \frac{1}{2} & \Longrightarrow & p(z_i = 1 \mid t_i = 1, R_i) = 1 \\ p_{R_i}(z_j = 1 \mid \theta = 1) > \frac{1}{2} & \Longrightarrow & p(z_i = 1 \mid t_i = 1, R_i) = 0 \end{array}$$

That is, trader i's decision whether to be at the location ψ when he learns it depends only on whether he finds it more likely that his opponent will also be at ψ when an opportunity arises. The claim allows us to define any symmetric equilibrium entirely in terms of the notation

$$\alpha_R = p(z = 1 \mid t = 1, R)$$

where equilibrium symmetry allows us to remove the subscript i from z, t, R.

Rational-expectations benchmark

Consider a benchmark model in which every realization of R induces rational expectations - i.e., $p_R(z = 1 \mid \theta = 1) = p(z = 1 \mid \theta = 1)$. The only possible difference between realizations of R is thus the value of q_R they induce. This reduces the model to a conventional game with incomplete information. While it is possible to sustain symmetric Nash equilibria in which traders' behavior is type-dependent, all these equilibria generate zero profits for both traders, independently of their type:

$$p(z_i = 1 \mid \theta = 1) = \frac{1}{2} \tag{5}$$

In particular, there is a symmetric Nash equilibrium in which traders' behavior is independent of their type:

$$\alpha_R = \frac{1}{2\sum_R p(R) q_R} \tag{6}$$

The zero-profit property conventionally captures the idea that market competition dissipates profit opportunities.

Let us now introduce a distribution over traders' R-information that allows for non-rational expectations. Assume trader i only receives archival data about the three variables θ, t_j, z_j (where $j \neq i$). These variables describe whether a trade opportunity arises, whether the opponent is informed of the location of trade opportunities and whether he chooses the right location. By Claim 1, these variables pin down trader i's payoff, hence the trader's subjective expected payoff will be well-defined even though he lacks archive-information about other variables. We can now explicitly define traders' R-information, which takes two possible values - given the shorthand notation 0 and 1:

$$\begin{aligned} R_i &= 0 : \{\{l_\theta, l_{t_j}\}, \{l_{t_j}, l_{z_j}\}\} \\ R_i &= 1 : \{\{l_\theta, l_{z_j}\}\} \end{aligned}$$

That is, $R_i = 1$ means that trader i learns the joint distribution of θ and z_j, whereas $R_i = 0$ means that trader i learns how each of these variables is jointly distributed with t_j. The probability of each of these values of R_i is $\frac{1}{2}$, independently of (θ, ψ).

When $R_i = 1$, trader i's beliefs are consistent with rational expectations: $p_{R_i=1}(z_j = 1 \mid \theta = 1) = p(z_j = 1 \mid \theta = 1)$. In contrast, when $R_i = 0$, trader

i's conditional prediction is

$$p_{R_i=0}(z_j = 1 \mid \theta = 1) = \sum_{t_j} p(t_j \mid \theta = 1) p(z_j = 1 \mid t_j) = p(t_j = 1 \mid \theta = 1) p(z_j = 1 \mid t_j = 1) \tag{7}$$

where the second equality follows from the observation (made in Claim 1) that $p(z_j = 1 \mid t_j = 0) = 0$.

The following elaboration of these formulas highlights the role of R_j as a confounder of the relation between t_j and z_j; $p_{R_i=1}$ properly accounts for this role,

$$p_{R_i=1}(z_j = 1 \mid \theta = 1) = \sum_{R_j} p(R_j) p(t_j = 1 \mid \theta = 1, R_j) p(z_j = 1 \mid t_j = 1, R_j)$$

whereas $p_{R_i=0}$ neglects it and wrongly presumes that $z_j \perp \theta \mid t_j$, such that $p_{R_i=0}(z_j = 1 \mid \theta = 1)$ can be written as

$$\left(\sum_{R_j} p(R_j) p(t_j = 1 \mid \theta = 1, R_j) \right) \left(\sum_{R_j} p(R_j \mid t_j = 1) p(z_j = 1 \mid t_j = 1, R_j) \right)$$

Importantly, if there were no variation in R_j, this confounding effect would disappear and $p_{R_i=0}(z_j = 1 \mid \theta = 1)$, too, would coincide with the rational-expectations prediction. That is, learning the marginals $p(\theta, t_j)$ and $p(t_j, z_j)$ would suffice for trader i to arrive at rational expectations. We are thus witnessing a similar "*bootstrapping*" effect to the one in Section 1.1. The dependence between t_j and z_j is confounded by the variation in trader j's type; and the data limitations of one of these types prevent him from detecting the same confounding effect that drives his opponent's behavior.

Observe that we can still sustain the type-independent strategy (6) in equilibrium. The reason is that under this strategy, traders' R-information ceases to function as a confounder, and therefore $p_{R_i=0}$ does not distort the objective mapping from θ to z_j. The question is whether there are symmetric equilibria that give rise to *heterogeneous* performance by market agents. The following result provides the answer.

Proposition 5 *(i) When $q_1 \geq q_0$, the only symmetric equilibrium is the one where traders play the type-independent strategy (6) and earn zero profits.*

(ii) When $q_1 < q_0$, there is exactly one additional symmetric equilibrium: A trader of type $R = 1$ plays $\alpha_1 = 1$ and earns a strictly positive profit, whereas a trader of type $R = 0$ plays $\alpha_0 = (2 - q_0 - q_1)/(q_0 + q_1)$ and earns zero profits.

Proof. Consider a symmetric equilibrium. By assumption, $p(z = 1 \mid t = 0, R) = 0$ for both $R = 0, 1$. It must be the case that $\alpha_0 > 0$ or $\alpha_1 > 0$ - otherwise, it would be profitable for any trader i to deviate to $z_i = 1$ when $t_i = 1$ because he would make a strictly positive profit when $\theta = 1$. Likewise, it must be the case that $\alpha_0 < 1$ or $\alpha_1 < 1$ - otherwise, traders earn negative expected payoff in equilibrium, whereas playing $z = 0$ would generate zero profit. As observed earlier, a trader of type R chooses $z = 1$ with positive probability only if $p_R(z = 1 \mid \theta = 1) \leq \frac{1}{2}$. Therefore, $\alpha_1 > \alpha_0$ only if

$$p_{R=1}(z = 1 \mid \theta = 1) < p_{R=0}(z = 1 \mid \theta = 1) \tag{8}$$

Let us derive explicit expressions for the two sides of this inequality:

$$p_{R=1}(z = 1 \mid \theta = 1) = p(z = 1 \mid \theta = 1) = \frac{1}{2}q_1\alpha_1 + \frac{1}{2}q_0\alpha_0$$

whereas

$$p_{R=0}(z = 1 \mid \theta = 1) = \left(\frac{1}{2}q_1 + \frac{1}{2}q_0\right)\left(\frac{1}{2}\alpha_1 + \frac{1}{2}\alpha_0\right)$$

The latter expression is obtained by plugging the terms

$$p(t_j = 1 \mid \theta = 1) = \frac{1}{2}q_1 + \frac{1}{2}q_0$$

and

$$\begin{aligned}
p(z_j &= 1 \mid t_j = 1) = \frac{p(t_j = z_j = 1)}{p(t_j = 1)} \\
&= \frac{\frac{1}{2}(\frac{1}{2}q_1\alpha_1 + \frac{1}{2}q_0\alpha_0) + \frac{1}{2}(\frac{1}{2}(1-q_1)\alpha_1 + \frac{1}{2}(1-q_0)\alpha_0)}{\frac{1}{2}} \\
&= \frac{1}{2}\alpha_1 + \frac{1}{2}\alpha_0
\end{aligned}$$

Then, (8) becomes

$$\frac{1}{2}q_1\alpha_1 + \frac{1}{2}q_0\alpha_0 < \left(\frac{1}{2}q_1 + \frac{1}{2}q_0\right)\left(\frac{1}{2}\alpha_1 + \frac{1}{2}\alpha_0\right)$$

which is equivalent to

$$\alpha_1(q_1 - q_0) < \alpha_0(q_1 - q_0)$$

Suppose $q_1 \geq q_0$. Then, this inequality contradicts the inequality $\alpha_1 > \alpha_0$. A similar contradiction is obtained for $\alpha_1 < \alpha_0$. It follows that when $q_1 \geq q_0$, it is impossible to sustain an equilibrium in which $p_{R=1}(z = 1 \mid \theta = 1) \neq p_{R=0}(z = 1 \mid \theta = 1)$ - hence, the only possible equilibrium is one where $\alpha_1 = \alpha_0 \in (0,1)$. Therefore, the only possible symmetric equilibrium is the one given by (6), such that traders earn zero profits. This establishes part (i).

Now suppose $q_1 < q_0$. Then, the above contradiction is not reached, and it is possible to sustain equilibria with $\alpha_1 \neq \alpha_0$. If $\alpha_1 \in (0,1)$, traders of type $R = 1$ are indifferent between $z = 0$ and $z = 1$. This indifference condition immediately gives (5), which means that all traders earn zero objective profits, independently of their type. Because $\alpha_1 \neq \alpha_0$, the subjective expected profit of traders of type $R = 0$ is not zero, hence $\alpha_0 \in \{0,1\}$, which contradicts (5). The only remaining possibility is $\alpha_1 = 1$ and $\alpha_0 \in (0,1)$, such that traders of type $R = 0$ are indifferent between $z = 0$ and $z = 1$. Plugging this indifference into $p_{R=0}(z = 1 \mid \theta = 1)$ yields the solution for α_0. Plugging this value in $p_{R=1}(z = 1 \mid \theta = 1)$, we can verify that traders of type $R = 1$ earn positive profits, consistent with the assumption that $\alpha_1 = 1$. ■

This result sheds light on the question of differential market performance among diversely sophisticated participants. In order to sustain equilibria in which some traders earn non-zero profits, we must have $q_1 < q_0$ - i.e., the two aspects of market savvy should be *negatively correlated*. If a trader with superior news-information also has better archive-information, then traders' behavior in symmetric equilibrium must be independent of their type and they all must earn zero profits; the sophisticated type's "double advantage" (superior news- and archive-information) has no effect on its market performance. The non-standard equilibrium that exists when $q_1 < q_0$ is unique: the sophisticated type $R = 1$ always chooses to be active when he is tipped off, whereas the "naive" type $R = 0$ only does so with some probability. Both types earn positive objective profits, yet the naive type overestimates the competition it will face and believes it makes zero profits; therefore, it does not become fully active.

The lesson is that the aspect of market savvy that matters for relative market performance is archive-information, which governs traders' understanding of other traders' equilibrium behavior. Recall that under rational expectations, superior news-information gives traders no benefit because the potential gains from spotting trade opportunities are dissipated in equilibrium. In contrast, differential archive-information means that some traders fail to correctly assess this "competitive dissipation". This in turn distorts their market behavior relative to the rational-expectations benchmark. If the direction of this distortion is downwards, trade opportunities are only partially dissipated and traders with superior archive-information have an edge. However, in order for this to happen, the correlation between traders' news- and archive-information must be negative.

The analysis highlights another feature of the model, which is the importance of the precise variables on which players have archival data. The coarse signal and consequence variables t_i and z_i are irrelevant under rational expectations; the model is completely defined by the variables $\theta, s_1, s_2, a_1, a_2$. However, in the present context the variables are important because they feed players' archival information. If we assumed that type $R_i = 0$ receives finer data and learns the marginal distributions $p(\theta, \psi, s_j)$ and $p(s_j, a_j)$, his equilibrium belief would be effectively reduced to rational expectations and all traders would earn zero profits in symmetric equilibrium.

5 Discussion of Related Literature

The literature contains a number of important game-theoretic solution concepts in which players receive partial feedback regarding equilibrium behavior, such that players' beliefs are based purely on this feedback (without any display of strategic introspection). It is helpful to define each of these proposals by two ingredients: the way it formalizes partial feedback and the belief-formation rule it assumes. The crucial novelty of this paper compared with existing approaches is that it includes an explicit model of players' uncertainty (including limited feedback) regarding other players' feedback.

The closest approaches to the one in this paper are those in which players extrapolate a belief from their feedback according to an explicit rule that intuitively follows a parsimony principle. Osborne and Rubinstein (1998) assume that a player's feedback takes the form of a collection of finite samples taken from the conditional distributions over outcomes that is induced by each action. Players ignore sampling error and believe that the sample associated with each action is perfectly representative of its true conditional distribution over outcomes. Salant and Cherry (2019) extend this basic idea to more general estimation procedures. Osborne and Rubinstein (2003) study a variant of this concept, in which each player's feedback consists of a sample drawn from the *unconditional* distribution over the opponent's actions. In

Esponda (2008), the feedback sample is infinite but *selective*. For example, in a bilateral trade example, it is the distribution of outcomes conditional on trade. Players' extrapolated belief reflects unawareness of the sample's selectiveness.

Jehiel (2005) and Jehiel and Koessler (2008) present a formalism that is closest in spirit to the present paper, in the sense that a player's feedback limitation is a *personal characteristic* rather than part of the definition of the solution concept. Under this approach, each player best-replies to a coarse representation of the true equilibrium distribution. Specifically, the player partitions the set of possible contingencies (histories in extensive games, states of the world in Bayesian games) into "analogy classes", such that the feedback that he receives is the average distribution over contingencies within each analogy class. His belief does not allow for finer variation within each analogy class. In Section 2 I showed how the present formalism can express this belief-formation model. Thus, at least in the context of Bayesian games, the archive-information formalism is a generalization of analogy-based expectations.

In other approaches that are closer to the tradition of self-confirming equilibrium (Battigalli (1987), Fudenberg and Levine (1993)), players do not extrapolate a belief from limited feedback. Instead, they arrive at the game with a subjective, possibly misspecified prior model, and they fit this model to their feedback. For example, Esponda and Pouzo (2016) formalize feedback abstractly as a general consequence variable (in applications, it typically coincides with the player's payoff, or with the realized terminal history in an extensive game). Each player has a prior belief over a set of possible distributions over consequences conditional on the game's primitives and the players' actions. This set represents the player's model, and it is misspecified if it rules out the true conditional distribution. In equilibrium, the player's belief is a conditional distribution in this set that is closest (according to a modified Kullback-Leibler Divergence) to the true equilibrium distribution.[7] Battigalli et al. (2015) assume a similar notion of feedback and adopt a non-probabilistic model of beliefs in the decision-theoretic "ambiguity" tradition.

These two general approaches to belief formation - extrapolation from feedback vs. fitting a subjective model to feedback - are not mutually exclusive. In particular, in Section 2 we saw how p_R can sometimes be interpreted as the result of fitting a subjective causal model to the objective distribution. Note, however, that here the player's feedback is *not* independent of his prior model. E.g., when the player's causal model is $\theta \rightarrow s \rightarrow a$, his feedback is the marginals $p(\theta, s)$ and $p(s, a)$.

Eyster and Rabin (2005) adopt a different interpretation of distorted equilibrium beliefs. In "fully cursed" equilibrium, player i wrongly believes that the distribution over a_j is a measurable function of player i's (!) signal. In "partially cursed" equilibrium, a player's belief is a convex combination between the rational-expectations and fully cursed beliefs. Eyster and Rabin regard this belief distortion as a behavioral bias and do not attempt to derive it from explicit partial feedback or from an explicit subjective model. However, one can easily reinterpret fully cursed beliefs along these lines (see Jehiel and Koessler (2008)). And Spiegler (2017a) provides a partial feedback-based justification for partially cursed beliefs.

A "self-confirming equilibrium" formulation

The formalism presented in this paper builds on a specific notion of players' feedback: learning marginal distributions over collections of variables. The following is a more abstract formulation, which is closer in spirit to the "self-confirming equilibrium" tradition.

Let Y be a set of observable outcomes. Define the state space as $X = \Theta \times S_1 \times S_2 \times F_1 \times F_2 \times A_1 \times A_2$, where F_i is a collection of "feedback functions" $f : X \rightarrow \Delta(Y)$. Player i's type is given by the pair (s_i, f_i) -

[7] We saw in Section 2 how the maximum-entropy extrapolation rule could be reformulated interms of minimizing Kullback-Leibler Divergence. The concept plays a very different role in Esponda and Pouzo (2016).

i.e., his news-information s_i and the feedback function f_i. Unlike the rest of this paper, here Y and f_i are abstract objects with no explicit structure. As before, the prior $p \in \Delta(X)$ satisfies the conditional-independence property $a_i \perp (\theta, s_j, f_j, a_j) \mid (s_i, f_i)$. Player i's strategy is given by the conditional distribution $(p(a_i \mid s_i, f_i))$. A prior p and a feedback function f induce the outcome distribution $\mu^{p,f} \in \Delta(Y)$, given by $\mu^{p,f}(y) = \sum_x p(x) f(y \mid x)$.

In a self-confirming ε-equilibrium, when player i's feedback function is f_i, he has a subjective unconditional belief $\tilde{p}^{f_i} \in \Delta(X)$ that is consistent with his feedback, in the sense that $\mu^{\tilde{p},f_i} = \mu^{p,f_i}$. He then conditions this belief on s_i, f_i; and $p(a_i \mid s_i, f_i) > \varepsilon$ if a_i best-replies to this conditional subjective belief. We can then refine this equilibrium concept by imposing the maximum-entropy criterion - i.e., player i's subjective unconditional belief $\tilde{p}^{f_i}$ maximizes entropy among all the distributions $q \in \Delta(X)$ for which $\mu^{q,f_i} = \mu^{p,f_i}$.

This formulation subsumes the framework of this paper as a special case. In particular, the feedback function f_i generalizes the notion of R-information. This extra-generality may be viewed as an advantage. Nevertheless, I preferred to focus on the more special case in this paper because its concreteness vastly facilitates the construction and interpretation of economic examples.

6 Conclusion

Previous attempts to develop equilibrium models with non-rational expectations treated belief distortions as an aspect of the solution concept or as a permanent fixture of individual players. The formalism presented in this paper enriches the scope of this literature by including individual players' limited feedback (called their R-information) in the description of their *type*. The formalism describes R-information in terms of the collections of *variables* about which the player receives feedback. Because players' R-information itself can be a random variable, this language enabled us to capture new and realistic kinds of "high-order" reasoning. It also enabled us to explore the economic implications of the correlation between this novel aspect of players' types and other, more conventional aspects (e.g. the accuracy of their news-information).

A natural next step is to extend the formalism to dynamic strategic interactions, where a move by one player at an early decision node can determine another player's archive-information at a later decision node. Eliaz et al. (2018) is a step in this direction. In that paper, we present a cheap-talk model in which the sender controls strategically the receiver's N-information as well as his R-information, and explore the implications of this novel feature on the sender's ability to persuade the receiver.

References

[1] Aumann, R. (1987), Correlated Equilibrium as an Expression of Bayesian Rationality, Econometrica 55, 1-18.

[2] Battigalli, P. (1987), Comportamento Razionale ed Equilibrio nei Giochi e nelle Situazioni Sociali, unpublished undergraduate dissertation, Bocconi University, Milano.

[3] Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni and M. Marinacci (2015), Self-Confirming Equilibrium and Model Uncertainty, American Economic Review 105, 646-677.

[4] Carlsson, H. and E. Van Damme (1993), Global Games and Equilibrium Selection, Econometrica 61, 989-1018.

[5] Cover, T. and J. Thomas (2006), Elements of Information Theory, second edition, Wiley.

[6] Cowell, R., P. Dawid, S. Lauritzen and D. Spiegelhalter (1999), *Probabilistic Networks and Expert Systems,* Springer, London.

[7] Eliaz, K., R. Spiegler and H. Thysen (2018), Strategic Interpretations, mimeo.

[8] Ellis, A. and M. Piccione (2017), Correlation Misperception in Choice, American Economic Review 107, 1264-92.

[9] Enke, B. and F. Zimmermann (2017), Correlation Neglecty in Belief Formation, Review of Economic Studies, forthcoming.

[10] Esponda, I. (2008), Behavioral Equilibrium in Economies with Adverse Selection, The American Economic Review 98, 1269-1291.

[11] Esponda. I. and D. Pouzo (2016), "Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models," Econometrica 84, 1093-1130.

[12] Ettinger, D. and P. Jehiel (2010), A Theory of Deception, American Economic Journal: Microeconomics 2, 1-20.

[13] Eyster, E. and M. Rabin (2005), Cursed Equilibrium, Econometrica 73, 1623-1672.

[14] Eyster, E. and M. Piccione (2013), An Approach to Asset Pricing Under Incomplete and Diverse Perceptions, *Econometrica,* 81, 1483-1506.

[15] Fudenberg, D. and D. Levine (1993), Self-Confirming Equilibrium, Econometrica 61, 523-545.

[16] Hajek, P., T. Havranek and R. Jirousek (1992), Uncertain Information Processing in Expert Systems, CRC Press.

[17] Hellwig, M. (2008), The Causes of the Financial Crisis, CESifo Forum 9 (4), 12-21.

[18] Jaynes, E. T. (1957), Information Theory and Statistical Mechanics, Physical Review 106, 620-630.

[19] Jehiel, P. (2005), Analogy-Based Expectation Equilibrium, Journal of Economic Theory 123, 81-104.

[20] Jehiel, P. (2018), Investment Strategy and Selection Bias: An Equilibrium Perspective on Overoptimism, American Economic Review 108, 1582-97.

[21] Jehiel, P. and F. Koessler (2008), Revisiting Games of Incomplete Information with Analogy-Based Expectations, Games and Economic Behavior 62, 533-557.

[22] Levy, G. and R. Razin (2015), Correlation Neglect, Voting Behaviour and Information Aggregation, American Economic Review 105, 1634-1645.

[23] Mailath, G. and L. Samuelson (2018), The Wisdom of a Confused Crowd: Model-Based Inference, mimeo.

[24] Morris, S. and H. Shin (2003), Global Games: Theory and Applications. In Advances in Economics and Econometrics (Proceedings of the Eighth World Congress of the Econometric Society), edited by M. Dewatripont, L. Hansen, and S. Turnovsky, 1:56-114. Cambridge University Press.

[25] Piccione, M. and A. Rubinstein (2003), Modeling the Economic Interaction of Agents with Diverse Abilities to Recognize Equilibrium Patterns, *Journal of the European Economic Association*, 1, 212-223.

[26] Osborne, M. and A. Rubinstein (1998), Games with Procedurally Rational Players, American Economic Review 88, 834-847.

[27] Osborne, M. and A. Rubinstein (2003), Sampling Equilibrium with an Application to Strategic Voting, Games and Economic Behavior 45, 434-441.

[28] Rubinstein, A. (1989), The Electronic Mail Game: A Game with Almost Common Knowledge, American Economic Review 79, 385-391.

[29] Salant, Y. and J. Cherry (2019), Statistical Inference in Games, mimeo.

[30] Selten, R. (1975), Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games, International Journal of Game Theory 4, 25-55.

[31] Spiegler, R. (2016), Bayesian Networks and Boundedly Rational Expectations, Quarterly Journal of Economics 131, 1243-1290.

[32] Spiegler, R. (2017), "Data Monkeys": A Procedural Model of Extrapolation From Partial Statistics, Review of Economic Studies 84, 1818-1841.

[33] Spiegler, R. (2018), Can Agents with Causal Misperceptions be Systematically Fooled? Journal of the European Economic Association, forthcoming.

[34] Woodford, M. (2013), "Macroeconomic Analysis without the Rational Expectations Hypothesis," *Annual Review of Economics* 5, 303-346.