

News and Archival Information in Games*

Ran Spiegler[†]

Previous version: March 2018

This version: August 2018

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. Building 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.” I demonstrate the expressive power and use of this formalism with some examples.

*Financial support by ERC Advanced Investigator grant no. 692995 is gratefully acknowledged. I thank Kfir Eliaz and Heidi Thysen 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

When people engage in a static strategic interaction, they make use of various kinds of information. One kind concerns the *current* realization of exogenous variables. Another kind consists of background data about joint realizations of exogenous and endogenous variables in past instances of the game. I refer to these two kinds as news-information (or *N*-information) and archive-information (or *R*-information), respectively. To use a journalistic metaphor, *N*-information is akin to a newsflash about a corporate scandal, whereas *R*-information is what a reporter gets when he starts digging the newspaper’s archives for evidence about the behavior of various actors in past scandals. While *N*-information is knowledge of characteristics of the current strategic situation, *R*-information enables the player to make sense of this knowledge and draw conclusions about the possible consequences of his actions.

Standard game theory treats these two types of information very differently. Harsanyi’s model of static games with incomplete information offers a rich general description of players’ incomplete information regarding the current realization of exogenous variables, including high-order information. In other words, Harsanyi’s type-space formalism is exclusively about news-information; it leaves the task of describing archive-information to the solution concept. Introspective, “one-shot” solution concepts like rationalizability or level- k reasoning ignore archive-information altogether. At the other extreme, Nash equilibrium presumes that players have complete archive-information.

The last two decades saw various proposals for solution concepts that retain the steady-state approach of Nash equilibrium, while relaxing its assumption of complete archive-information and replacing it with some notion of limited learning feedback that players receive regarding the steady-state distribution, coupled with some model of how players form beliefs given their partial feedback. I provide a detailed literature review in Section 5. At this stage, it suffices to say that virtually all previous proposals assume that the

feedback limitation that characterizes each player is fixed. And neither provides a model of players' imperfect information (of either kind) regarding their opponents' archive-information.

However, it is easy to think of real-life situations in which both types of information fluctuate. For instance, a sophisticated player may be characterized by accurate news-information as well as rich archival information, whereas a naive player will be deprived of both. It is also natural to think of situations in which one player has incomplete news-information about 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 behavior in past situations, some of which share the present situation's exogenous characteristics. Likewise, we can meaningfully talk about one player having incomplete archive-information about another player's archive-information. For instance, army 1 may receive access to a computer file that documents army 2's archival access in other situations. And one can easily extend these descriptions to high-order statements that involve both news-information and archive-information, just as we do for news-information in the standard Harsanyi model.

In this paper I present a new type space for static games, which combines both kinds of incomplete information. 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, the game's outcome). 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 has 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 partially informed about the current realization of the exogenous variables. The novelty lies in the *second component* of the player’s type, namely his *R*-information. This is defined as a *collection of subsets* of 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 subsets of variables defined by his *R*-information. Thus, rather than learning the entire joint distribution p , 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 over the variables about which he has data, thus forming a potentially distorted perception of the objective distribution p . In the second stage, the player conditions this extrapolated belief on his type (both his *N*-information and his *R*-information) via Bayes’ rule, to form a subjective belief over payoff-relevant outcomes as a function of his action. Equilibrium is defined in a standard way: essentially, each player plays a best-reply to his subjective belief.

Of course, there are many extrapolation rules one could employ in the procedure’s first stage. However, a recurring theme in the literature 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 do not have 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

I capture this motive by assuming that players use the *maximum-entropy*

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

extrapolation criterion. That is, the player’s belief is the distribution (over the variables about which he has archival data) that maximizes entropy subject to being consistent with the marginals he learns. This extrapolation is “parsimonious” in the sense that it does not postulate correlations that lack basis in his data. As I later show, it subsumes existing notions in the literature as special cases, and easy to calculate and visualize in many applications of interest.

Thus, each component of the player’s type corresponds to a different stage in his belief-formation process. The first stage makes use of the player’s R -information via the extrapolation rule of maximum entropy, while the second stage makes use of his N -information via the conditioning rule of Bayesian updating. To revisit the journalistic metaphor, players first use “archival research” to extrapolate an unconditional belief, and then condition it on the “news flash”.

1.1 An Example: 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’s order of moves or the opponent’s preferences and mode of reasoning. Each player interacts *once*, after getting some feedback about the behavior of previous generations agents who assumed the players’ current roles. If players knew that 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 in the game is based *entirely* on their (possibly incomplete) learning feedback, which is 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. A state of the world is described by the quadruple (R_1, R_2, a_1, a_2) . 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 any steady-state distribution over the four variables. With probability α , players have incomplete archive-information:

$$\begin{aligned} R_1 &= \{\{l_{R_1}, l_{a_1}\}, \{l_{a_1}, l_{a_2}\}\} \\ R_2 &= \{\{l_{R_2}, l_{a_2}\}, \{l_{a_1}, l_{a_2}\}\} \end{aligned}$$

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. Players do not receive any news-information. In particular, player i does not receive any signal regarding R_j . Thus, player i 's type is defined solely by R_i .

The interpretation is as follows. The distribution over (a_1, a_2) represents a large, publicly available record of past game realizations that were drawn from the steady-state joint distribution p over all four variables. Because the record is public, players always have access to it and they can learn the steady-state distribution over the action profile. In contrast, the record of past joint realizations of player i 's archive-information and his action need not be public. Indeed, it is privately accessed by player i alone with probability α . With probability $1 - \alpha$, the complete historical record of all variables becomes publicly available.

The exogenous component of the prior p is $p(R_1, R_2)$ - i.e., the distribution over players' types. The endogenous components are the players' strategies, given by the conditional probability distributions $p(a_1 | R_1)$ and $p(a_2 | R_2)$. Because the game has simultaneous/independent moves, p satisfies the conditional independence properties $a_1 \perp (R_2, a_2) | R_1$ and $a_2 \perp (R_1, a_1) | R_2$.

When player i 's type is R_i , he forms his belief in two stages. First, he extrapolates an unconditional subjective belief p_{R_i} over the variables he has data on. Second, he conditions this belief on his type and action to evaluate the action's payoff consequences. Let us derive players' beliefs as a function of their types. Complete archive-information means rational expectations. That is, when $R_i = \{L\}$, $p_{R_i} = p$. Because D is a strictly dominant action, it follows that when player i 's type is $R_i = \{L\}$, he will necessarily play D in any equilibrium. In contrast, suppose that $R_i = \{\{l_{R_i}, l_{a_i}\}, \{l_{a_1}, l_{a_2}\}\}$. Then, the player learns the marginal distributions $p(R_i, a_i)$ and $p(a_1, a_2)$. The maximum-entropy extension of these marginals is

$$p_{R_i}(R_i, a_1, a_2) = p(R_i, a_i)p(a_j | a_i)$$

Conditioning this belief on the player's news information $t_i = R_i$ and his action, we obtain

$$p_{R_i}(a_j | R_i, a_i) = p(a_j | a_i)$$

Thus, when player i 's type is $R_i = \{\{l_{R_i}, l_{a_i}\}, \{l_{a_1}, l_{a_2}\}\}$, he forms a conditional subjective belief regarding a_j as if he thinks that his own action causes the opponent's action. In other words, he acts as if he mistakes the correlation between a_i and a_j (due to their respective dependence on players' correlated types) for a *causal effect* of the former on the latter.

This crucial feature of the player's belief is a direct consequence of the maximum-entropy extension criterion - which in turn relies on the principle that players are reluctant to postulate correlations beyond those for which they have evidence for. Furthermore, players do not use any other criterion

for constructing beliefs about their opponents' behavior - a plausible assumption when players lack an understanding of the game's structure. It makes sense for a player who is unaware that he is playing a Prisoner's Dilemma - as opposed to, say, a Stackelberg-like game in which the opponent has a coordination motive - to form a belief that his action directly causes the opponent's action, because this is a parsimonious interpretation of the correlations he learns.

Equilibrium in this environment is defined conventionally, in the spirit of trembling-hand perfection (Selten (1975)). A profile of completely mixed strategies constitutes an ε -equilibrium if whenever $p(a_i | R_i) > \varepsilon$, a_i maximizes player i 's expected utility with respect to his conditional subjective belief $p_{R_i}(\cdot | R_i, a_i)$. An equilibrium is simply a limit of a sequence of ε -equilibria, where $\varepsilon \rightarrow 0$. We can now characterize the set of symmetric equilibria in this example. When equilibria are not sensitive to the perturbation's form, I will omit this part and go straight to the equilibria, without describing ε -equilibria first.

One equilibrium is for players to play D regardless of their type - this is the conventional game-theoretic prediction. Because C is never played, this equilibrium needs to be sustained by a perturbation. Suppose that players choose D with probability $1 - \varepsilon$, independently of their type, where ε is arbitrarily small. Then, $p(a_j | a_i) = 1 - \varepsilon$ for all a_i , and therefore $a_i = D$ is strictly dominant, hence the players' strategies constitute an ε -equilibrium. Taking the $\varepsilon \rightarrow 0$ limit gives us the equilibrium.

Another equilibrium is for each player i to play D if and only if $R_i = \{L\}$. To see why, we only need to establish that when $R_i = \{\{l_{R_i}, l_{a_i}\}, \{l_{a_1}, l_{a_2}\}\}$, $a_i = C$ is a best-reply (because D is strictly dominant when $R_i = \{L\}$). Under this candidate equilibrium, $p(a_j = C | a_i = C) = p(a_j = D | a_i = D) = 1$. We saw that when $R_i = \{\{l_{R_i}, l_{a_i}\}, \{l_{a_1}, l_{a_2}\}\}$, player i 's belief can be described as if he interprets the perfect correlation between a_i and a_j causally, and therefore believes that player j will play C if and only if i plays

C . Therefore, 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 $R_i = \{\{l_{R_i}, l_{a_i}\}, \{l_{a_1}, l_{a_2}\}\}$. (As before, he plays D with probability one when $R_i = \{L\}$.) Best-replying requires the player to be indifferent between the two actions, given his conditional subjective belief:

$$\begin{aligned} 3 \cdot p(a_j = C \mid a_i = C) + 0 \cdot p(a_j = D \mid a_i = C) \\ = 4 \cdot p(a_j = C \mid a_i = D) + 1 \cdot p(a_j = D \mid a_i = D) \end{aligned}$$

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

$$\lambda = \frac{1}{3 - 2\alpha}$$

Thus, the formalism enables us to capture the idea that some types of players perceive any long-run correlation between players’ actions in causal terms. The common variation in players’ archive-information creates the correlation between their actions, and one realization of their type gives rise to the mistaken causal interpretation of this correlation. Of course, correlated actions could arise from other sources, e.g. correlated payoff shocks. What is novel about the present example is that the fundamental correlation is in players’ *non-payoff* types, and one of these types lacks access to archival data that correctly accounts for this correlation. This sort of “bootstrapping” effect will be a running theme in this paper. Novel effects will occur because players’ archive-information is random, and because some realizations of this random variable lead to distorted perception of the correlation patterns that arise from this source of randomness.

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 $u_i : X \rightarrow \mathbb{R}$. Let $p \in \Delta(X)$ be an objective prior distribution over the state space.

To reflect the simultaneous-move assumption, p satisfies the following conditional independence properties: $a_i \perp (\theta, s_j, R_j, a_j) \mid (s_i, R_i)$. That is, player i 's action is independent of the other exogenous variables and player j 's action conditional on his type. 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 $\sigma_2 = (p(a_2 \mid s_2, R_2))$. Unless indicated otherwise, p has full support.

Comment: The notion of a state

A state of the world resolves *all* uncertainty, including the 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 players, and p records the collective experience of many other agents who assumed the two 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 4.3 we will see 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 $x = (x_1, \dots, x_n)$, 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 transparently associate a variable with its label. Let $L = \{1, \dots, n\}$ 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 the prior p on X_B .

Now, any realization of R_i is a particular *collection of subsets of the set of variable labels* 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 restrict R_i to satisfy the property that u_i is measurable with respect to $x_{N(R_i)}$. That is, the player never lacks archival data about a variable that is necessary to define his payoffs unambiguously. I

also typically assume that $l_{a_i} \in N(R_i)$ - i.e., player i always obtains archival data regarding a_i . The latter assumption is not necessary and only made to simplify notation at certain points.

Belief formation

To make a decision, player i forms a probabilistic belief as a function of his type. In the standard Harsanyi model, 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 that makes use of different kinds 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

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

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.²

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

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

²The solution to the constrained maximization problem will always be unique.

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. In particular, a number of solution concepts in the literature (see Section 5) involve parsimonious treatment of *correlations*. When players lack sufficient understanding of the game’s structure or their opponent’s preferences, this is a plausible criterion for extrapolating beliefs from partial learning feedback.

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

$$\sum_{z \in X_{N(R_i)}} q(z) \frac{\ln(q(z))}{\ln(q^*(z))} \quad (2)$$

This formulation is equivalent because q^* is uniform and therefore $\ln(q^*(z))$ is a constant. Expression (2) 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.

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 $p(a_i | s_i, R_i) > \varepsilon$,

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

A strategy profile (σ_1^*, σ_2^*) (which need not satisfy full support) is an equilibrium if it is the limit of a sequence of ε -equilibria with $\varepsilon \rightarrow 0$.

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.

2.1 More on Maximum-Entropy Extrapolation

Calculating the maximum-entropy extension of a given collection of marginal distributions can be non-trivial. However, in many cases it takes a tractable, interpretable form. Indeed, some of these special cases have effectively appeared in the literature on equilibrium models with non-rational expectations; the maximum-entropy criterion can be regarded as a principle that unifies them.

The simplest case is where $R_i = \{B\}$, where $B \subset L$. That is, player i receives complete data about the joint distribution of a particular subset of variables. In this case, $p_{R_i} = p^B$ - this is trivially pinned down by the requirement that the player's belief is consistent with the marginal distribution he has learned; there is no need to apply maximum entropy. While this

may seem like a trivial case, it can have significant behavioral implications, arising from the fact that the player’s belief is defined over X_B rather than X . E.g., suppose that the realization $R_i = \{\{l_{R_i}, l_{a_i}\}, \{l_{a_1}, l_{a_2}\}\}$ in Section 1.1 is replaced with $\{\{l_{a_1}, l_{a_2}\}\}$. Then, player i ’s belief over a_j conditional on (R_i, a_i) remains the same, and therefore the equilibrium analysis is unchanged. 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)).

Another simple example is when R consists of mutually disjoint subsets of L . The maximum-entropy extension in this case is

$$p_R = \prod_{B \in R} p^B$$

E.g., consider the realization $R_1 = \{\{l_\theta\}, \{l_{a_2}\}\}$, which only records the marginal distributions $(p(\theta))$ and $(p(a_2))$ without conveying any data about the correlation between the two variables. Then, p_{R_1} is the distribution q over (θ, a_2) that minimizes

$$\sum_{\theta, a_2} q(\theta, a_2) \ln q(\theta, a_2)$$

subject to the constraints

$$\sum_{a_2} q(\theta, a_2) = p(\theta) \text{ for all } \theta \quad \text{and} \quad \sum_{\theta} q(\theta, a_2) = p(a_2) \text{ for all } a_2$$

We can now write down the first-order conditions of the Lagrangian of this constrained minimization problem and obtain the solution $p_{R_1}(\theta, a_2) \equiv p(\theta)p(a_2)$. Thus, the maximum-entropy extension of the marginals over individual variables treats them as mutually independent.

Slightly more complicated examples involve R -information realizations that consist of intersecting subsets. E.g., suppose that $x = (\theta^1, \theta^2, R_1, R_2, a_1, a_2)$, where θ^1 and θ^2 are two components that define the state of Nature. Consider

the realization $R_1 = \{\{l_{\theta^1}, l_{\theta^2}\}, \{l_{\theta^2}, l_{a_2}\}\}$, indicating 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

$$p_R(\theta^1, \theta^2, a_2) = p(\theta^1, \theta^2)p(a_2 | \theta^2)$$

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.

Mailath and Samuelson (2018) analyze information aggregation under a similar model of belief formation. 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 (which in principle can be identified with their actions). Mailath and Samuelson assume that the player's subjective belief over these variables is $p(b)p(x_M | b)p(z | x_M)$. This is the maximum-entropy extension of the marginals $p(b, x_M)$ and $p(x_M, z)$, and therefore can be described in terms of the present formalism.

The preceding examples are all special cases of a more general result, which I briefly present now as all the examples in this paper utilize it.³

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

RIP holds trivially for $m = 2$. The $m = 3$ collection $\{\{1, 2\}, \{2, 3\}, \{3, 4\}\}$ satisfies the property, whereas the $m = 3$ collection $\{1, 2\}, \{2, 3\}, \{1, 3, 4\}$

³The material here is taken from Spiegler (2017a), where I studied a behaviorally motivated procedure for extrapolating beliefs from partial data. This procedure coincides with maximum entropy when the following combinatorial property is satisfied.

violates it. Hajek et al. (1992) show that when R satisfies RIP, the maximum-entropy extension of $(p^B)_{B \in R}$ is given by

$$p_R(x_{N(R)}) = \prod_{B^1, \dots, B^m} p(x_{B^k - (\cup_{j < k} B^j)} \mid x_{B^k \cap (\cup_{j < k} B^j)}) \quad (3)$$

where the enumeration $1, \dots, m$ validates the running intersection property. For instance, when $R = \{\{1, 2\}, \{2, 3\}, \{3, 4\}\}$,

$$\begin{aligned} p_R(x_1, x_2, x_3, x_4) &= p(x_1, x_2)p(x_3 \mid x_2)p(x_4 \mid x_3) \\ &= p(x_1)p(x_2 \mid x_1)p(x_3 \mid x_2)p(x_4 \mid x_3) \end{aligned} \quad (4)$$

Thus, RIP allows p_R to be written as a factorization of $p(x_{N(R)})$ into marginal and conditional distributions. Moreover, the factorization has a *causal* interpretation. For instance, (4) looks as if it is consistent with the causal chain $1 \rightarrow 2 \rightarrow 3 \rightarrow 4$.

This is a general property. Suppose that R satisfies RIP. Define a *directed acyclic graph* $G = (N, E)$, where $N = N(R)$ is the set of nodes and $E \subset N \times N$ is the set of directed links (that is, $(j, i) \in E$ means that there is a link $j \rightarrow i$). Slightly abusing notation, define $G(i) = \{j \in N \mid (j, i) \in E\}$. A subset of nodes $C \subset N$ is a *clique* in G if there is a link between any pair of nodes in C (that is, for every $i, j \in C$, $i \neq j$, $(i, j) \in E$ or $(j, i) \in E$). A clique is *maximal* if it is not contained in another clique. For any prior p , define p_G as the factorization of p according to G , given by the formula

$$p_G(x_N) = \prod_{i \in N} p(x_i \mid x_{G(i)})$$

This is known in the literature as a Bayesian-network factorization formula (e.g., see Cowell et al. (1999)). The following is a restatement of a result by Hajek et al. (1992) - see Spiegler (2017a) for more details.

Proposition 1 *Suppose that R satisfies RIP. Then, there exists a directed acyclic graph G such that $p_G \equiv p_R$. Moreover, G satisfies two additional*

properties. First, it is perfect - i.e., $G(i)$ is a clique for every $i \in N$. Second, the set of maximal cliques in G is R .

This result establishes that when R satisfies RIP, the maximum-entropy extension of the agent’s archival data can be equivalently described as the outcome of fitting a *subjective causal model* - formalized as a directed acyclic graph - to the objective distribution p . Moreover, R is the set of maximal cliques in the graph. Finally, the graph is perfect, which implies a number of important and useful properties (see Spiegler (2017a,b)). In particular, it induces *correct marginal distributions* over all individual variables.

The graphical representation of p_R under RIP has two important roles. First, it offers a convenient visualization of the player’s archive-information. This in turn makes the conditional independence properties of p_R visible and facilitates calculations. Second, it highlights the feature that the player’s extrapolated belief looks as if he imposes a false causal interpretation on observed correlations - as we saw in Section 1.1. In all the examples presented in this paper, players’ R -information satisfies RIP.

3 Correlation between News and Archival Information: A Market Competition Game

In this section and the next I offer a number of examples that illustrate the formalism. In each case, I show how the formalism can efficiently capture a realistic, hitherto unmodeled feature of real-life strategic interactions. One such feature is correlation between the two components of a player’s type - namely, his news-information and archive-information. In some real-life situations, players who are more knowledgeable about the current situation are generally “savvier” and therefore also know more about the game’s history. In other cases, we may view the two components as substitutes that

arise from an unmodeled prior stage of information acquisition: A player can become an expert on current affairs or on history but not on both.

The following example develops this theme in the context of a stylized market competition game. The state of the world is $x = (\theta, s_1, s_2, R_1, R_2, a_1, a_2)$, where $\theta \in \{0, 1\}$ indicates whether there is demand for a certain product. The players are firms whose market access is uncertain: $s_i = 1$ (0) means that the technology for making the product is available (unavailable) to firm i . When $s_i = 0$, the firm's only feasible action is $a_i = 0$ (inactivity). In contrast, when $s_i = 1$, the firm can also choose to be active in the market (an action denoted $a_i = 1$). Note that firm i 's action set varies with its news-information s_i . Firm i 's payoff function is $u_i(\theta, a_1, a_2) = \theta a_i(d - a_j)$, where $d \in (0, \frac{1}{2})$ represents the size of market demand. The interpretation is that a single active firm is a monopolist that earns positive profits, whereas two active firms lead to competitive pressures that result in negative profits.

The firm's R -information can take two values, referred to by the shorthand notation 0 and 1 and given explicitly as follows:

$$\begin{aligned} R_i = 0 & : \{ \{l_\theta, l_{s_1}\}, \{l_\theta, l_{s_2}\}, \{l_{s_1}, l_{a_1}\}, \{l_{s_2}, l_{a_2}\} \} \\ R_i = 1 & : \{ \{l_\theta, l_{s_1}, l_{a_1}\}, \{l_\theta, l_{s_2}, l_{a_2}\} \} \end{aligned}$$

The exogenous components of the prior p are as follows:

- (i) $p(\theta = 1) = \frac{1}{2}$.
- (ii) $p(R_i = 1) = \frac{1}{2}$, independently of θ and of R_j .
- (iii) For every $i = 1, 2$ and every θ , $p(s_i = \theta \mid \theta, R_i) = q_{R_i}$ independently of R_j , where $q_1, q_0 \geq \frac{1}{2}$. Denote

$$k = \frac{2d}{q_0 + q_1} \in (0, 1)$$

This parameter measures the size of market demand relative to firms' expected market access.

Let us interpret the notion of a player's type in this example. The variable

s_i is a conventional signal that determines firm i 's market access. When it is highly correlated with θ , the firm is “savvy”: it tends to have the technology for delivering a product in demand. However, “savviness” has another dimension, given by the firm’s archive-information. When $R_i = 1$, firm i fully grasps the joint distribution of demand and any individual firm’s market access and behavior. In contrast, $R_i = 0$ means that the firm lacks direct evidence regarding this distribution: it only learns the pairwise correlations of individual firms’ market access with their behavior and with market demand. The parameters q_1 and q_0 determine the correlation between firms’ N -information and R -information. When $q_1 > q_0$ ($q_1 < q_0$) the correlation is positive (negative): a firm with richer archive-information is more (less) likely to have a technology that matches demand.

The assumption that firms earn zero payoffs for sure when $\theta = 0$ implies that firms’ beliefs regarding θ are irrelevant: the only thing that matters for firm i 's decision is its prediction of a_j conditional on $\theta = 1$. In particular, firm i chooses a_i to maximize

$$a_i \cdot [d - p_{R_i}(a_j = 1 \mid \theta = 1)]$$

Both realizations of R_i satisfy RIP. Furthermore, it is easy to see, using the tools of Section 2.1, that p_{R_i} treats a_j, s_j as independent of a_i, s_i conditional on θ . Since the only thing that matters for player i 's belief is the mapping from θ to a_j , we can simplify the definition of R_i w.l.o.g:

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

That is, the only relevant aspect of a firm’s archive-information is the data about the joint distribution of θ and *its competitor’s* market access and behavior. When $R_i = 1$, firm i 's archive-information fully documents the joint distribution of θ and a_j . Therefore, its conditional prediction is consistent

with rational-expectations:

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

In contrast, when $R_i = 0$, firm i 's conditional prediction is

$$p_{R_i=0}(a_j = 1 \mid \theta = 1) = \sum_{s_i} p(s_j \mid \theta = 1) p(a_j = 1 \mid s_j)$$

By the assumption that firm j is forced to play $a_j = 0$ when $s_j = 0$, this can be simplified into

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

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

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

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

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

This effect can also be illustrated graphically, using the tools of Section 2.1 and focusing on the relevant variables θ, s_j, a_j . The belief $p_{R_i=0}(a_j = 1 \mid \theta = 1)$ is consistent with the DAG $\theta \rightarrow s_j \rightarrow a_j$, whereas the belief $p_{R_i=1}(a_j = 1 \mid \theta = 1)$ is consistent with a DAG that adds a direct link $\theta \rightarrow a_j$ to this graph.

Importantly, if firms' R -information were constant, this confounding ef-

fect would disappear and $p_{R_i=0}(a_j = 1 \mid \theta = 1)$, too, would coincide with the rational-expectations prediction. We are thus witnessing a similar “*bootstrapping*” effect to the one observed in Section 1.1: the dependence between firms’ market access and market behavior is confounded by their type variation; and the data limitations of one of these types prevent him from detecting this confounding effect.

The following analysis focuses on symmetric equilibria. I refer to $p(a = 1 \mid \theta = 1)$ as the *market-activity rate* that characterizes the equilibrium. I use the notation

$$\alpha_R = p(a = 1 \mid s = 1, R)$$

to describe the equilibrium strategy of firms of type R .

Rational-expectations benchmark

Because the precision of firms’ signals is immaterial for their decisions, the example has a stark rational-expectations benchmark. Suppose we assumed that the only difference between the realizations $R_i = 0$ and $R_i = 1$ lies in the precision of the signal s_i - i.e., both types have rational expectations, and their only difference is that $q_1 \neq q_0$. Then, we would have a conventional game with incomplete information, in which each firm i receives a conditionally independent signal s_i with type-dependent accuracy. Because both firm types in this alternative model have rational expectations, they make the same prediction of a_j conditional on $\theta = 1$. While it is possible to sustain symmetric Nash equilibria in which firms’ mixture over actions (conditional on having market access) is type-dependent, all these equilibria have the same market-activity rate

$$p(a = 1 \mid \theta = 1) = d \tag{5}$$

In particular, there is a symmetric Nash equilibrium in which

$$\alpha_1 = \alpha_0 = k \tag{6}$$

Let us now turn to the original specification of the example, where $R_i = 0$ potentially induces wrong beliefs. Observe that we can sustain an equilibrium in which firms play the type-independent strategy (6), such that the market-activity rate is given by (5). The reason is that under such a strategy profile, firms' R -information ceases to function as a confounder, and therefore $p_{R_i=0}$ does not distort the objective mapping from θ to a_j . The question is whether there are symmetric equilibria that give rise to *different* market-activity rates. The following result provides the answer.

Proposition 2 (i) *When $q_1 \geq q_0$, the only symmetric equilibrium is the one where firms play the type-independent strategy (6). As a result, the equilibrium market-activity rate coincides with the rational-expectations benchmark (5).*

(ii) *When $q_1 < q_0$, there is one symmetric equilibrium in which the market-activity rate departs from the rational-expectations benchmark. In particular:*

(iia) *If $k > \frac{1}{2}$, then $\alpha_0 = 2k - 1$, $\alpha_1 = 1$, and the market-activity rate is $\frac{1}{2}[q_1 + q_0(2k - 1)] < d$.*

(iib) *If $k < \frac{1}{2}$ then $\alpha_0 = 2k$, $\alpha_1 = 0$, and the market-activity rate is $q_0k > d$.*

Proof. Consider a symmetric equilibrium. By assumption, $p(a = 1 \mid s = 0, R) = 0$ for both $R = 0, 1$. Recall the notation $\alpha_R = p(a = 1 \mid s = 1, R)$. Because $d < 1$, it must be the case that in any symmetric equilibrium, $\alpha_0 > 0$ or $\alpha_1 > 0$, as well as $\alpha_0 < 1$ or $\alpha_1 < 1$. As observed earlier, a firm of type R chooses to be active ($a = 1$) with positive probability only if $p_R(a = 1 \mid \theta = 1) \leq d$. Therefore, $\alpha_1 > \alpha_0$ only if

$$p_{R=1}(a = 1 \mid \theta = 1) < p_{R=0}(a = 1 \mid \theta = 1) \quad (7)$$

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

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

whereas

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

The latter expression is obtained by plugging the terms

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

and

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

Then, (7) 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 $p_{R=1}(a = 1 \mid \theta = 1) \neq p_{R=0}(a = 1 \mid \theta = 1)$ - 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 the market-activity rate is (5). 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 $\alpha_1 \in (0, 1)$, firms of type $R = 1$ are indifferent between the actions. But their indifference condition immediately gives (5). Therefore, in order to sustain an equilibrium with a different market-activity rate, we need $\alpha_1 = 0$ or $\alpha_1 = 1$, such that $\alpha_0 \in (0, 1)$ - i.e., firms of type $R = 0$ are indifferent between the two actions. Plugging each of these cases $\alpha_1 = 1$ or $\alpha_1 = 0$ into $p_{R=0}(a = 1 | \theta = 1)$ yields the two equilibria described in part (ii). ■

This result shows how the framework constrains our ability to capture differential market performance among “diversely sophisticated” firms. In order to sustain equilibria in which the market-activity rate departs from the rational-expectations benchmark, it must be the case that $q_1 < q_0$ - i.e., there is *negative correlation* between the two “savviness” components. If a firm with superior news-information also has better archive-information, firms’ behavior in symmetric equilibrium must be independent of their type, and therefore their performance is also the same: the sophisticated type’s “double advantage” (superior news- and archive-information) has no effect on its market outcome!

The non-standard equilibrium that exists when $q_1 < q_0$ is unique, yet its precise structure depends on the size of market demand d . When d is relatively small, the “sophisticated” firms of type $R = 1$ choose to be inactive, while the “naive” firms of type $R = 0$ are active with positive probability. The latter firms perceive the market-activity rate to be equal to d and predict zero profits. The actual market-activity rate is above d , such that active firms make negative profits. In contrast, when d is relatively high, firms of type $R = 1$ choose to be active, while the “naive” firms of type $R = 0$ are inactive with positive probability. As in the previous case, the latter firms perceive the market-activity rate to be equal to d and predict zero profits. The actual

market-activity rate is below d , such that being active brings positive profits.

4 Information about Archive-Information: A Coordination Game

In this section I illustrate another aspect of the formalism’s expressive scope - specifically, how it enables us to incorporate novel, realistic aspects of high-order reasoning. 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 payoff structure makes high-order strategic reasoning crucial for players’ behavior, and therefore enables us to illustrate the novel types of high-order reasoning that the present 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.

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

$$\begin{array}{rcc}
 a_1 \backslash a_2 & 1 & 0 \\
 1 & \delta\theta, \delta\theta & -1, 0 \\
 0 & 0, -1 & 0, 0
 \end{array}$$

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 $(1, 1)$.

4.1 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 $\alpha \in (0, \frac{1}{2})$, $R_2 = \{\{l_\theta, l_{a_1}, l_{a_2}\}\}$ - a realization also denoted $R_2 = 1$. 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) beliefs 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 the first two special cases presented in Section 2.1. 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 act considerately in a social situation.

When $\theta = 0$, players cannot afford to be considerate. When $\theta = 1$, there are gains from mutually considerate 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 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 3 *There is a unique equilibrium in this example: Both players always play $a = 0$.*

Proof. 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 $\theta = 0$.

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 $a_1 = 1$.

4.2 Archive-Information about Archive-Information

In this sub-section, I analyze a situation in which equilibrium patterns arise from players' partial 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 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). The DAG representation of $R_i = 0$ is $a_j \leftarrow \theta \rightarrow a_i$, whereas the DAG representation of $R_i = 1$ is $R_i \leftarrow R_j \leftarrow a_j \leftarrow \theta \rightarrow a_i$.

The causal interpretation of these DAGs is absurd because it treats R_j as a consequence of a_j . Nevertheless, it facilitates the derivation of players' conditional beliefs. The DAGs make it clear that p_{R_i} satisfies the conditional-independence property $a_j, R \perp a_i \mid \theta$. Therefore, we can ignore a_i when calculating player i 's belief over θ, a_j :

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

In this example (as in Section 3), 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 $R_i = 0$ become de-facto identical, and therefore there is no reason for player i to vary his action with R_i . And as in Section 3, when we look for equilibria that allow players' actions to vary with their archive-information, only particular configurations can be sustained.

Proposition 4 *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

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

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 $p_{R_i=1}(a_j = 1 \mid \theta = 1, R_i = 1)$, by plugging the following terms

$$\begin{aligned} p(a_j = 1 \mid \theta = 1) &= \frac{1}{2} \\ p(R_j = 1 \mid a_j = 1) &= 1 \end{aligned}$$

and

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

Thus,

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

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 a symmetric pure-strategy equilibrium in which $a_i = \theta R_i$ for every player i . Any other equilibrium in this category 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, such that it induces rational expectations. Then, the above equilibrium can be sustained for *any* $\delta > 0$.

4.3 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 the standard Harsanyi model, collapsing the collections θ , s_i or z into a single variable can entail a loss of generality. For instance, suppose that the state of Nature θ has multiple dimensions, such that $\theta = (\theta^1, \dots, \theta^K)$. Player i 's archive-information R_i may only provide partial data about the joint distribution of these components - i.e., for every $k = 1, \dots, K$ there is $B \in R_i$ that includes l_{θ^k} , yet there exists no $B \in R_i$ that contains $\{l_{\theta^1}, \dots, l_{\theta^K}\}$. The belief distortions that arise from this partial R -information cannot be reproduced if we collapse θ into a single variable.

Perhaps the most interesting case of this proliferation of variables is where *players' archive-information itself* is represented by a collection of variables. In the examples of Sections 1.1 and 4.2, 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 Harsanyi formalism's ability to define one player's N -information regarding another player's N -information. And as in the Harsanyi model, 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 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, \dots, R_i^m , $m \geq 2$, where R_i^1 is a collection of subsets of B ; and for every $k = 2, \dots, m$, R_i^k is a collection of subsets of $B \cup \{l_{R_i^1}, \dots, l_{R_i^{k-1}}\} \cup \{l_{R_j^1}, \dots, l_{R_j^{k-1}}\}$, where each of these subsets includes $l_{R_i^{k-1}}$ or $l_{R_j^{k-1}}$. Define $R_i = \cup_{k=1, \dots, m} R_i^k$.

The interpretation of this hierarchical construction is natural: 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 our coordination game. 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 $k = 1, 2, \dots$, R^k takes two values, 0 and 1, which are 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	$\{\{l_\theta, l_{a_1}, l_{a_2}, l_{R^1}, \dots, l_{R^{m-1}}\}\}$
\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 $p(R = k) = \delta(1 - \delta)^k$ for every $k = 0, 1, \dots$. Note that $R = k$ means that players perceive actions as a function of $\theta, R^1, \dots, R^{k-1}$.

Proposition 5 *Suppose that $\delta > \frac{1}{2}$. Then, there is a unique equilibrium, in which players always play $a = 0$.*

Proof. The proof is by induction on R . 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 $p_{R=k}(a_2 = 1 \mid \theta = 1, R = k, a_1 = 1) > \frac{1}{2}$,

but this does not hold since

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

■

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 $\delta > \frac{1}{2}$, 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.

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, and where players' beliefs are purely based on partial learning, 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 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 the correlation-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. Osborne and Rubinstein (2003) study a

variant of this concept, in which each player’s feedback consists of one sample drawn from the *unconditional* distribution over outcomes. In Esponda (2008), the feedback sample is not finite but it is *selective* - i.e., it is drawn from the distribution over outcomes conditional on players’ equilibrium behavior. Players’ rule for extrapolating from the sample reflects unawareness of its 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-responds to the following coarsening of the true equilibrium distribution: 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 2.1 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 (Fudenberg and Levine (1993)), players do not extrapolate a belief from limited feedback. Instead, they arrived at the game with a subjective, possibly misspecified prior model, and they fit this model to the 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.⁴ 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 2.1 we saw how p_R can sometimes be interpreted as the result of fitting a prior 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 DAG is $\theta \rightarrow s \rightarrow a$, his feedback enables him to learn the joint distributions over θ, s and s, a .

Eyster and Rabin (2005) take a different interpretation to modeling distorted equilibrium beliefs in games. In “fully cursed” equilibrium, a player wrongly believes that the distribution over his opponents’ actions is a measurable function of his own 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)). Likewise, I will show that the “model-based inference” in Mailath and Samuelson (2018) can be described by the language of the present formalism.

The crucial difference between all the approaches described above and the present paper is that in all these papers, limited feedback is either part of the solution concept or a fixed, non-random characteristic of each player. (The only exception I am familiar with is the Appendix to Eyster and Rabin (2005), where the parameter that defines the degree of a player’s cursedness is drawn from some distribution.) In contrast, in the present paper players’

⁴We saw in Section 2 how the maximum-entropy extrapolation rule could be reformulated in terms of minimizing Kullback-Leibler Divergence. However, the concept plays a very different role in Esponda and Pouzo (2016).

limited feedback is a random variable that constitutes an aspect of his type. And none of the existing approaches involves an explicit model of players' uncertainty (including limited feedback) regarding other players' feedback.

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, and subsumes the present formalism as a special case. I briefly describe this alternative formalism, in order to clarify how the present formalism relates to that 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) - 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-independent property that a_i is independent of θ, s_j, f_j, a_j conditional on 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 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 ; and $p(a_i \mid s_i, f_i) > 0$ if a_i best-responds to this conditional subjective belief. We can 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. Never-

theless, I preferred to focus on the more special case in this paper because its concreteness facilitated the construction of examples, and because its structure enabled the link to the idea that players form beliefs as if they fit a subjective causal model to objective data. The more general version is considerably more abstract and harder to make use of when trying to write down economic applications - a challenge that I leave for future research.

6 Conclusion

This paper took the idea of equilibrium modeling without rational expectations as a point of departure. Although the literature has explored a variety of modeling approaches to this general idea, virtually all attempts treat the departure from rational expectations 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 equilibrium modeling with non-rational expectations, 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. This language enables us to capture new and realistic kinds of “high-order” reasoning, such as *N*-information or *R*-information about another player's *R*-information, as well as hierarchical constructions of higher order. 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.

References

- [1] Aumann, R. (1987), Correlated Equilibrium as an Expression of Bayesian Rationality, *Econometrica* 55, 1-18.
- [2] Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni and M. Marinacci (2015), Self-Confirming Equilibrium and Model Uncertainty, *American Economic Review* 105, 646-677.
- [3] Carlsson, H. and E. Van Damme (1993), Global Games and Equilibrium Selection, *Econometrica* 61, 989-1018.
- [4] Cover, T. and J. Thomas (2006), *Elements of Information Theory*, second edition, Wiley.
- [5] Cowell, R., P. Dawid, S. Lauritzen and D. Spiegelhalter (1999), *Probabilistic Networks and Expert Systems*, Springer, London.
- [6] Eliaz, K., R. Spiegler and H. Thyssen (2018), Strategic Interpretations, in progress.
- [7] Enke, B. and F. Zimmermann (2017), Correlation Neglect in Belief Formation, *Review of Economic Studies*, forthcoming.
- [8] Esponda, I. (2008), Behavioral Equilibrium in Economies with Adverse Selection, *The American Economic Review* 98, 1269-1291.
- [9] Esponda, I. and D. Pouzo (2016), "Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models," *Econometrica* 84, 1093-1130.
- [10] Ettinger, D. and P. Jehiel (2010), A Theory of Deception, *American Economic Journal: Microeconomics* 2, 1-20.
- [11] Eyster, E. and M. Rabin (2005), Cursed Equilibrium, *Econometrica* 73, 1623-1672.

- [12] Eyster, E. and M. Piccione (2013), An Approach to Asset Pricing Under Incomplete and Diverse Perceptions, *Econometrica*, 81, 1483-1506.
- [13] Fudenberg, D. and D. Levine (1993), Self-Confirming Equilibrium, *Econometrica* 61, 523-545.
- [14] Hajek, P., T. Havranek and R. Jirousek (1992), Uncertain Information Processing in Expert Systems, CRC Press.
- [15] Hellwig, M. (2008), The Causes of the Financial Crisis, CESifo Forum 9 (4), 12-21.
- [16] Jaynes, E. T. (1957), Information Theory and Statistical Mechanics, *Physical Review* 106, 620-630.
- [17] Jehiel, P. (2005), Analogy-Based Expectation Equilibrium, *Journal of Economic Theory* 123, 81-104.
- [18] Jehiel, P. and F. Koessler (2008), Revisiting Games of Incomplete Information with Analogy-Based Expectations, *Games and Economic Behavior* 62, 533-557.
- [19] Levy, G. and R. Razin (2015), Correlation Neglect, Voting Behaviour and Information Aggregation, *American Economic Review* 105, 1634-1645.
- [20] Mailath, G. and L. Samuelson (2018), The Wisdom of a Confused Crowd: Model-Based Inference, mimeo.
- [21] 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.

- [22] 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.
- [23] Osborne, M. and A. Rubinstein (1998), Games with Procedurally Rational Players, *American Economic Review* 88, 834-847.
- [24] Osborne, M. and A. Rubinstein (2003), Sampling Equilibrium with an Application to Strategic Voting, *Games and Economic Behavior* 45, 434-441.
- [25] Rubinstein, A. (1989), The Electronic Mail Game: A Game with Almost Common Knowledge, *American Economic Review* 79, 385-391.
- [26] Selten, R. (1975), Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games, *International Journal of Game Theory* 4, 25-55.
- [27] Spiegler, R. (2016), Bayesian Networks and Boundedly Rational Expectations, *Quarterly Journal of Economics* 131, 1243-1290.
- [28] Spiegler, R. (2017a), "Data Monkeys": A Procedural Model of Extrapolation From Partial Statistics, *Review of Economic Studies* 84, 1818-1841.
- [29] Spiegler, R. (2017b), Can Agents with Causal Misperceptions be Systematically Fooled? CEPR Discussion Paper 11379.
- [30] Woodford, M. (2013), "Macroeconomic Analysis without the Rational Expectations Hypothesis," *Annual Review of Economics* 5, 303-346.