News and Archival Information in Games*

Ran Spiegler†

March 28, 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 correlations among relevant variables. A player is characterized by the conventional Harsanyi type (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. The player extrapolates a well-specified probabilistic belief 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 news flash 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. 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 1.2. 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 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 types of incomplete information. This formalism adds a bit of pedantry to the standard model, by requiring the game's description to include an explicit list of *variables* and classify them as exogenous variables, players' actions and consequence variables. It also makes heavy use of an explicit notational distinction between variables and their labels. Denote the set of variable labels by L. A player's type has two components. The first component is the player's N-information, defined as the *realization* of some subcollection of the exogenous variables. This is essentially the standard, Harsanyi notion of a player's type: 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", to use a different metaphor. For any possible steady-state probability distribution p over all variables, the player gets to learn the marginal of p over each of the subsets of variables that correspond to the subsets of L constituting the player's R-information. In other words, rather than learning the entire joint distribution p, the player only has piecemeal knowledge of it, given by some marginals. Although the player may fail to receive archival data about certain variables, I assume that his payoff function is always measurable with respect to the variables about which he does get data.

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 distorted perception of the objective steady-state distribution p. In the second stage, the player conditions this extrapolated belief on his N-information via Bayes' rule, to form a subjective belief over payoff-relevant outcomes as a function of his action. Equilibrium is defined in a completely standard way: each player always plays a best-reply to his subjective belief.

Of course, there are infinitely many extrapolation rules one could follow in the procedure's first stage. However, a recurring theme in the literature is that players apply an "Occam's razor" parsimony criterion when extrapolating from limited feedback. In this spirit, I assume that players use the *maximum-entropy* 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 knows. This extrapolation is "parsimonious" in the sense that it does not postulate correlations that lack a basis in his data.

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.

Let R_i and a_i denote player i's R-information and action, respectively. The set of variable labels is $L = \{\delta_{a_1}, \delta_{a_2}, \delta_{R_1}, \delta_{R_2}\}$. With probability $1-\alpha$, both players have complete archive-information - that is, $R_1 = R_2 = \{L\}$. This means that the players have a full grasp any steady-state distribution over all four variables. With probability α, players have incomplete archive-information, where R_1 and R_2 are given as follows:

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

That is, player i learns the joint distribution over his archive-information and action, as well as the joint distribution over the action profile.

The interpretation is as follows. The joint distribution over (a_1, a_2) represents a large, publicly available record of past game outcomes, each independently drawn from a 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 may be privately accessed by player i alone. Indeed, this is the case with probability α. With probability $1-\alpha$, however, the complete historical record of all variables becomes publicly available.

Finally, let us define the players' conventional news-information. Assume $t_i = R_i$ - i.e., each player knows what archival data he has access to, but he does not know his opponent's. This assumption is in fact immaterial for the equilibrium analysis of this example, but it simplifies the exposition because it means that player i's type is entirely pinned down by his R-information.

The exogenous component of 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 $p(a_1 \mid R_1)$ and $p(a_2 \mid R_2)$. The assumption of simultaneous/independent moves means that p satisfies the conditional-independence properties $a_1 \perp R_2, a_2 \mid R_1$ and $a_2 \perp R_1, a_1 \mid R_2$. 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. Then, he conditions this belief on his news-information t_i and his action a_i to evaluate the payoff consequences of his actions.

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 $R_i = \{\{\delta_{R_i}, \delta_{a_i}\}, \{\delta_{a_1}, \delta_{a_2}\}\}$. 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

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

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

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

Thus, when player i's type is $R_i = \{\{\delta_{R_i}, \delta_{a_i}\}, \{\delta_{a_1}, \delta_{a_2}\}\}$, 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.

Equilibrium in this environment is defined conventionally, in the spirit of trermbling-hand perfection (Selten (1975)). A profile of completely mixed strategies constitutes an ε-equilibrium if whenever $p(a_i \mid t_i, R_i) > \varepsilon$, a_i maximizes player i's expected utility with respect to his conditional subjective belief $p_{R_i}(\cdot \mid t_i, a_i)$. An equilibrium is simply a limit of a sequence of ε-equilibria, where $\varepsilon \to 0$. 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, $p(a_j \mid a_i) = 1-\varepsilon$ for all a_i, and therefore $a_i = D$ is strictly dominant, hence the players' strategies constitute an ε-equilibrium. Taking the $\varepsilon \to 0$ 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 $R_i = \{\{\delta_{R_i}, \delta_{a_i}\}, \{\delta_{a_1}, \delta_{a_2}\}\}$, $a_i = C$ is a best-reply (because D is strictly dominant under rational expectations). Under this candidate equilibrium, $p(a_j = C \mid a_i = C) = p(a_j = D \mid a_i = D) = 1$. We saw that when $R_i = \{\{\delta_{R_i}, \delta_{a_i}\}, \{\delta_{a_1}, \delta_{a_2}\}\}$, 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 $R_i = \{\{\delta_{R_i}, \delta_{a_i}\}, \{\delta_{a_1}, \delta_{a_2}\}\}$. (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:

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

Plug

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

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 archive-information gives rise to the mistaken causal interpretation of this correlation.

1.2 Related Literature

As mentioned earlier in the Introduction, the literature contains a number of important game-theoretic solution concepts in which players receive partial feedback regarding equilibrium behavior. 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 approaches in the literature that are closest to the present paper are those in which players' partial feedback takes a concrete form, and they extrapolate a belief from their feedback according to an explicit rule that intuitively follows the Occam's razor 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 the 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 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 I show that this belief-formation model can be expressed in the language of the present paper. Thus, at least in the context of Bayesian games, the archive-information formalism is a generalization of analogy-based expectations.

In other approaches, 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 (in the tradition of self-confirming equilibrium (Fudenberg and Levine (1993)) as a general consequence variable (in applications, it typically coincides with the player's payoff, or with the observed 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. 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 3 I will describe a link between the approach in this paper and Spiegler (2016), where decision makers fit a prior subjective (causal) model to feedback.

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)). Spiegler (2017a) shows that in some examples, partially cursed beliefs can be reconciled with a model of naive extrapolation from limited feedback that overlaps the maximum-entropy rule.

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 characteristic of each player. The modeler allows for no uncertainty about this feature of the game. 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. Neither of the existing approaches involves an explicit model of players' uncertainty regarding other players' feedback.

2 The Formalism

For expositional simplicity, I restrict attention to interactions that involve two players, 1 and 2. The extension to more than two players is entirely straightforward. Let $x = (x_1, ..., x_n)$ be a collection of variables, where $x_k \in X_k$ and X_k is some finite set. The *label* of variable x_i is i. When the notation for x_i has a particular meaning (as in the Introduction) I use the alternative notation l_{x_i}. Denote the set of variable labels by L. We can decompose x as follows, $x = (x_E, a_1, a_2, x_C)$, where E and C are two disjoint subsets of L, such that x_E is the collection of *exogenous* variables, a_1 and a_2 denote the players' actions, and x_C is the collection of *consequence* variables (in many applications, C will be empty). Player i's *type* is defined by a pair (R_i, t_i), where R_i denotes the player's archive-information, and t_i denotes his news-information. Both components are (possibly overlapping) projections of x_E. Each player i is a subjective expected utility maximizer, with vNM function $u_i : X \to \mathbb{R}$.

Let $p \in \Delta(X)$ represent a steady-state distribution over all variables. The most convenient interpretation of p is that it represents a historical record of many similar interactions. The individual game is a one-shot interaction between players who may lack direct, personal experience with similar situations - p records the collective experience of many other agents who assumed the roles of the two players.

The exogenous components of p are $(p(x_E))$ and $(p(x_C \mid x_E, a_1, a_2))$. Assume that these components have full support (in applications, we can sometimes afford to relax this restriction). The endogenous components of p are the players' strategies, $\sigma_i = (p(a_i \mid R_i, t_i))$ for every $i = 1, 2$. Strategies have full support, unless indicated otherwise. To reflect the assumption that players' moves are independent, assume that conditionally on t_i and R_i, a_i is independent of all other exogenous variables as well as of a_j.

Let us now impose explicit structure on R_i, in order to substantiate its interpretation as archive-information. Define R_i as a collection of subsets of L - i.e., the set of variable labels. For any steady-state distribution p and every $S \in R_i$, player i learns the marginal of p over x_S (denoted p^S). That is, R_i represents the player's limited access to archival data - his "database privileges", so to speak. 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 $(x_{N(R_i)}, a_i)$. That is, the player always has data on variables that unambiguously determine his payoffs. Note that I do not require $l_{a_i} \in N(R_i)$ - i.e., player i's archival information need not include data about a_i. On one hand, such data may be easier to retrieve (as explained in the Prisoner's Dilemma example). On the other hand, player i is mostly interested in predicting *player j's* behavior; learning about the past actions of agents who assumed the role of player i need not be relevant for this purpose.

To make a decision, player i forms a probabilistic belief over payoff-relevant variables, as a function of his type. In the standard model of games of incomplete information, the player's type is defined solely in terms of his news-information, and he forms his conditional probabilistic assessment in a single step, using Bayes' rule. In the present model, the 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.

The first stage involves *extrapolation from archive-information*: the player forms the unconditional belief $p_{R_i} \in \Delta(X_{N(R_i)})$ according to the *maximum-entropy* criterion. That is:

$$p_{R_i} = \arg \max_{q \in \Delta(X_{N(R_i)})} - \sum_{z \in X_{N(R_i)}} q(z) \ln(q(z))$$

subject to the constraint that $q \equiv p^S$ for every $S \in R_i$.[1] That is, the player's unconditional belief (over the variables about which he has data) maximizes entropy subject to being consistent with the marginals he learns. The second stage involves *conditioning on news-information*. Here the conditioning rule is conventional Bayesian updating. The player's conditional belief is $p_{R_i}(x_{N(R_i)} \mid t_i)$. This is the belief that the player uses to evaluate actions.

Thus, each component of the player's type is associated with a particular *operation* that he performs on the objective joint distribution. 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 should condition 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 an "Occam's razor" parsimony criterion. This principle is implicit in many of the solution concepts cited in the Introduction. The maximum-entropy criterion (which has a rich tradition in statistical physics and data analysis (see Jaynes (1957)) is one way to implement the parsimony criterion. It basically views minimal assumptions on correlations as parsimonious, and thus looks for the distribution that exhibits maximal independence, subject to being consistent with the correlation patterns that *are* observed in the data.

For illustration, let $L = \{1, 2\}$ and suppose that $R = \{\{1\}, \{2\}\}$. Suppose that x_i gets n_i values, such that any pair (x_1, x_2) can be identified with a distinct pair (i, j), $i = 1, ..., n_1$ and $j = 1, ..., n_2$. The known marginals are, for each i, j, $r_i = \sum_{j=1}^{n_2} p_{ij}$ and $s_j = \sum_{i=1}^{n_1} p_{ij}$. The maximum-entropy extension of these marginals is the joint probability distribution $(q_{ij})_{i,j=1,\ldots,n}$ that minimizes $\sum_{i=1}^{n_1} \sum_{j=1}^{n_2} q_{ij} \ln q_{ij}$ subject to the constraints $r_i = \sum_{j=1}^{n_2} q_{ij}$ and $s_j = \sum_{i=1}^{n_1} q_{ij}$. We can now write down the first-order conditions of the Lagrangian and obtain the solution $q_{ij} = r_i s_j$. That is, $p_R(x_1, x_2) =$

[1] The solution to the constrained maximization problem will always be unique.

$p(x_1)p(x_2)$.

Another example shows how to express analogy-based expectations (Jehiel (2005)) in terms of this formalism. Let $L = \{1, 2, 3\}$, where 1 is the label of a state of Nature, 2 is the label of an opponent's action, and 3 is the label of a variable that records the analogy class to which each state of Nature belongs. Suppose that $R = \{1, 3\}, (2, 3)\}$. Then, the player learns the marginals $p(x_1, x_3)$ and $p(x_2, x_3)$. The maximum-entropy extension of these marginals is $p(x_1)p(x_3 \mid x_1)p(x_2 \mid x_3)$, which is exactly what analogy-based expectations would prescribe.

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 steady-state distribution 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, t_i, R_i for which $p(a_i \mid t_i, R_i) > \varepsilon$,*

$$a_i \in \arg\max_{a_i'} \sum_{z \in X_{N(R_i)}} p_{R_i}(z \mid t_i, a_i')u_i(z, a_i')$$

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-equilibrium with $\varepsilon \to 0$.[2]

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.

I conclude this section with two comments about the formalism. First, the language I use to describe the strategic interaction is slightly different from the conventional description of simultaneous-move games. Instead of a game form, I represent the simultaneity assumption by a conditional independence

[2] The conditioning of p_{R_i} on a_i' is irrelevant if $l_{a_i} \notin N(R_i)$.

assumption that I impose on the joint distribution over all variables. In addition, the primitives of the model contain an explicit list of variables. In particular, the state of Nature or players' news-information or archive-information can be described as collections of variables. This level of detail is important. For example, when the state of Nature is described by multiple variables, the player need not have perfect archive-information about the correlation among these variables.

Second, the definition of the player's subjective belief $p_{R_i}(\cdot \mid t_i, a_i)$ does not make any explicit assumption about whether the player actively conditions on his own archive-information. One could assume that he is perfectly aware of his current archival-data limitations, and therefore conditions his belief on it. Alternatively, one might want to capture situations in which the player is not perfectly introspective, and therefore fails to realize that his current data limitations may have informational content. Any assumption we wish to make in this regard is fully encoded in the joint distribution of t_i and R_i. In particular, full introspection will be captured by a specification in which R_i is a component of t_i (as in the Prisoner's Dilemma example of the Introduction).

3 A Graphical Representation

Calculating the maximum-entropy extension of a collection of marginal distributions can be non-trivial. However, there are cases (like the ones I used to illustrate this rule in Section 2) in which it takes a simple form that is easy to interpret and work with. To take an extreme case, if all subsets that constitute R are disjoint, then

$$p_R(x_{N(R)}) = \prod\nolimits_{S \in R} p(x_S)$$

More generally, when R satisfies a certain combinatorial property, p_R can be described as if it is the result of fitting a *causal model* to the objective distribution p. In this section I briefly explain this observation. The material here is unoriginal, and borrowed from the literature on Bayesian networks (e.g., Cowell et al. (1999), Pearl (2009), Koller and Friedman (2009)).[3]

The collection R of subsets of L satisfies the *running intersection property* if the subsets that constitute R can be ordered $S^1, ..., S^m$ such that for every $k = 2, ..., m$, $S^k \cap (\cup_{j<k} S^j) \subseteq S^i$ for some $i = 1, ..., k-1$. This property obviously always holds for $m = 2$. The collection $\{\{1,2\},\{2,3\},\{3,4\}\}$ satisfies the property, whereas the collection $\{1,2\},\{2,3\},\{1,3,4\}\}$ violates it. It can be shown (see Hajek et al. (1992)) that when R satisfies the running intersection property, the maximum-entropy extension of $(p^S)_{S \in R}$ is given by

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

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

$$\begin{aligned} p_R(x_1, x_2, x_3, x_4) &= p(x_1, x_2) p(x_3 \mid x_2) p(x_4 \mid x_3) \\ &= p(x_1) p(x_2 \mid x_1) p(x_3 \mid x_2) p(x_4 \mid x_3) \end{aligned} \tag{2}$$

Thus, the running intersection property implies that p_R can be presented as a factorization of $p(x_{N(R)})$ into a product of marginal and conditional distributions. Moreover, the factorization has a *causal* interpretation. For instance, (2) looks as if it is consistent with a causal chain that can be described diagrammatically by the directed graph $1 \rightarrow 2 \rightarrow 3 \rightarrow 4$.

This is a general property. Suppose that R satisfies the running intersection property. 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

[3] The exposition is essentially taken from Spiegler (2017a), which examines a "behavioral" extrapolation procedure that coincides with maximum entropy when the combinatorial property is satisfied.

$G(i) = \{j \in N \mid (j,i) \in E\}$. As usual, 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 objective distribution 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.

The following is a restatement of a result by Hajek et al. (1992) - see Spiegler (2017a) for more details.

Proposition 2 *Suppose that R satisfies the running intersection property. 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 the running intersection property, 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. Moreover, the relation between the graph and R is simply that 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 interpretation of p_R has two important roles. First, it offers a convenient visualization of the player's type. This in turn makes the conditional independence properties of p_R visible and facilitates calculations. Second, it highlights the feature that the belief that the player extrapolates from partial data looks as if he imposes a false causal interpretation on observed correlations. Recall that in the Prisoner's Dilemma example of the

Introduction, a player's subjective belief could sometimes be described as if he falsely believes that his action causes the opponent's action and explained the observed long-run correlation between the players' actions in terms of this wrong causal perception.

It is important to note that the perfect-DAG representation of p_R is not unique. If two perfect DAGs R and R' have the same set of maximal cliques, then $p_R = p_{R'}$. For example, (2) could be equivalently written as

$$p_R(x_1, x_2, x_3, x_4) = p(x_4)p(x_3 \mid x_4)p(x_2 \mid x_3)p(x_1 \mid x_2)$$

which factorizes p according to the DAG $4 \rightarrow 3 \rightarrow 2 \rightarrow 1$. See Spiegler (2017a,b) for extended discussions of this property.

4 Examples

In this section I illustrate the formalism with three examples. My objective is to demonstrate the formalism's expressive scope - specifically, how it enables us to incorporate realistic aspects of high-order reasoning that are new to the literature on equilibrium models with non-rational expectations. 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.

Thus, throughout the section, we will 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	$\theta - \varepsilon, \theta - \varepsilon$	$-1, 0$
0	$0, -1$	$0, 0$

where $\varepsilon > 0$ is arbitrarily small, $\theta \in \{0, 1\}$ is the state of Nature, and $p(\theta = 1) = \frac{1}{2}$. In all the examples, players know the realization of θ. Therefore, under rational expectations, they would 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 $(1, 1)$.

In most of this section, I will assume that player i does not obtain archival data about his own past actions. That is, $l_{a_i} \notin N(R_i)$ for every realization of R_i. This means that the player's extrapolated belief treats a_j as independent of a_i conditional on player i's news-information. As a result, he will play $a_i = 0$ when $\theta = 0$, regardless of his R-information.

4.1 N-Information about R-Information

Suppose that player 1 has complete R-information with probability one. Therefore, we can omit R_1 as a variable. The only relevant variables are θ, R_2, a_1, a_2, where only θ, a_1, a_2 are directly payoff-relevant. Player 2's R-information is random and independent of θ. With probability $\alpha \in (0, \frac{1}{2})$, $R_2 = \{\{l_\theta, l_{a_1}\}\}$ - a realization also denoted $R_2 = 1$. With the remaining probability $1 - \alpha$, $R_2 = \{\{l_\theta\}, \{l_{a_1}\}\}$ - a realization also denoted $R_2 = 0$. Thus, R_2 records whether player 2 learns the correlation between player 1's action and the state of Nature.

As to players' N-information, assume for now that $t_1 = \theta$ and $t_2 = (\theta, R_2)$. That is, player 1 knows the current state of Nature but receives no information about player 2's current R-information - we will later relax this assumption. Player 2 knows the current state of Nature as well as his own current R-information.

Let us construct player 2's beliefs as a function of his type:

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

The derivation of $p_{R_2=1}$ is trivial, because $R_2 = 1$ enables player 2 to learn precisely the joint distribution of θ and a_1. The derivation of $p_{R_2=0}$ is essentially the same as in the example we used in Section 2 to illustrate the maximum-entropy rule.

Therefore,

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

Thus, the realization $R_2 = 0$ captures a "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 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 the interpretation of players' dilemma in this example is that they need to decide whether to act considerately in some social situation. In state $\theta = 0$, players cannot afford to be considerate, regardless of what the opponent does. In state $\theta = 1$, there are gains from mutually considerate behavior, and players are willing to act considerately if they expect similar behavior from the opponent. 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.

Proposition 3 *There is a unique equilibrium in this example, in which both players always play $a = 0$.*

Proof. When $\theta = 0$, both players choose $a = 0$. To see why, note first that player 1 and 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. Therefore, he finds $a_2 = 0$ to be the unique best-reply. The same argument holds when $R_2 = 0$ and $\theta = 1$.

Let us try to sustain an equilibrium in which $p(a_1 = 1 \mid \theta = 1) > 0$. As we saw in the previous paragraph, player 2 plays $a_2 = 0$ when $R_2 = 0$, independently of θ. Thus, player 2 plays $a_2 = 1$ with positive probability only if $\theta R_2 = 1$, such that $p(a_2 = 1 \mid \theta = 1) \leq \alpha$. Player 1's unique best-reply is $a_1 = 0$, a contradiction. It follows that player 1 always plays $a_1 = 0$ in any equilibrium. Completing the proof is straightforward. ■

The plot thickens: A second-order attribution error

What would happen if player 1 had partial N-information regarding player 2's R-information - i.e., he received a signal about whether player 2 exhibits the Fundamental Attribution Error? To incorporate this possibility, let us add the explicit variable $s_1 \in \{0, 1\}$, which represents player 1's signal regarding the value of R_2. Assume that the signal is binary with accuracy q - i.e., $p(s_1 = R_2) = q$ for every R_2, where $q \in (\frac{1}{2}, 1)$. Player 1' N-information is now $t_1 = (\theta, s_1)$. Player 2's N-information is as before: $t_2 = (\theta, R_2)$.

Continue to assume that player 1 has complete R-information. As to player 2, we now need to assume whether he has archive-information about the new variable s_1. Suppose that R_2 is distributed exactly as in the original specification of this example. In particular, player 2 never receives any data about s_1. Therefore, p_{R_2} is exactly as in the original specification. The difference is that now $R_2 = 1$ does *not* induce rational expectations. Rather, it captures a *"second-order" attribution error*: the player does not understand that player 1's behavior is responsive to his news-information about whether player 2 exhibits an attribution error.

Proposition 4 *There is a unique equilibrium in this example, in which both players always play $a = 0$.*

Proof. Let us try again 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:

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

Because p_{R_2} is as in the original specification, the result that player 2 plays

$a_2 = 1$ with positive probability only if $\theta R_2 = 1$ continues to hold. By our assumptions on α and q, $p(R_2 = 1 \mid s_1 = 0) < \frac{1}{2}$. Therefore, when player 1 observes $s_1 = 0$, his unique best-reply is $a_1 = 0$. It follows that

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

Therefore, 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 straightforward. ■

Note that player 2's second-order attribution error is key to this negative result. If $R_2 = 1$ represented complete R-information, player 2 would be able to infer from $R_2 = 1$ that player 1 is likely to observe $s_1 = 1$ and (if q is high enough) 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 does not condition his forecast of player 1's action on R_2 and ends up underestimating the conditional probability that player 1 will choose $a_1 = 1$.

4.2 R-Information about R-Information

Suppose that for each player i, $R_i = 0$ signifies $\{\{l_{R_1}, l_{R_2}\}, \{l_{R_j}, l_{a_j}\}, \{l_\theta\}\}$ and $R_i = 1$ signifies $\{\{l_{R_1}, l_{R_2}\}, \{l_{R_j}, l_{a_j}\}, \{l_\theta, l_{a_j}\}\}$. That is, player i does not always learn the joint distribution over (θ, a_j), but he always learns the joint distributions over the pairs of variables (R_1, R_2) and (R_j, a_j). Crucially, the player never learns how player j's action is correlated with the *joint* realization of player j's news-information and archive-information.

Assume that the R-information profile (R_1, R_2) is distributed independently of θ. The probability of (R_1, R_2) is denoted $\alpha_{R_1 R_2}$. The distribution is symmetric in the sense that $\alpha_{01} = \alpha_{10}$. Denote $\alpha = \alpha_{11} + \alpha_{10}$ (this is the probability that $R_i = 1$ for each i). Denote $q = \alpha_{11}/\alpha$ and assume $q < \frac{1}{2}$. This is the probability that $R_j = 1$ conditional on $R_i = 1$. Player i's N-information is $t_i = (\theta, R_i)$.

To derive player i's belief as a function of his type, we can employ the graphical representation presented in Section 3. The collection of subsets given by $R_i = 1$ satisfies the running intersection property. Therefore, we can view the three sets as maximal cliques in a perfect DAG, e.g.

$$R_i \to R_j \to a_j \to \theta \tag{3}$$

such that $p_{R_i=1}$ factorizes p according to this DAG:

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

Likewise, the three subsets that constitutes $R_i = 0$ satisfies the running intersection property, and can be viewed as maximal cliques in a perfect DAG that simply omits the link $a_j \to \theta$ from the DAG given by (3). Therefore,

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

It follows that

$$p_{R_i=0}(a_j \mid \theta, R_i, a_i) = \sum_{R_j} p(R_j \mid R_i)p(a_j \mid R_j)$$

$$p_{R_i=1}(a_j \mid \theta, R_i, a_i) = \frac{\sum_{R_j} p(R_j \mid R_i)p(a_j \mid R_j)p(\theta \mid a_j)}{\sum_{R_j} p(R_j \mid R_i) \sum_{a'_j} p(a'_j \mid R_j)p(\theta \mid a'_j)}$$

As in the previous example, and using essentially the same arguments, each player i necessarily plays $a_i = 0$ in equilibrium whenever $\theta R_i = 0$.

Proposition 5 *There is an equilibrium in which each player i plays $a_i = 1$ if and only if $\theta R_i = 1$, as long as*

$$q > \frac{2 - 2\alpha}{4 - 3\alpha}$$

Proof. Let us try to sustain this equilibrium. Best-replying for player i requires that

$$p_{R_i=1}(a_j = 1 \mid \theta = R_i = 1) > \frac{1}{2}$$

Given the players' assumed strategies,

$$\begin{aligned} p_{R_i=1}(a_j &= 1 \mid \theta = R_i = 1) = \frac{\sum_{R_j} p(R_j \mid R_i = 1) p(a_j = 1 \mid R_j) p(\theta = 1 \mid a_j = 1)}{\sum_{R_j} p(R_j \mid R_i = 1) \sum_{a'_j} p(a'_j \mid R_j) p(\theta = 1 \mid a'_j)} \\ &= \frac{q \cdot \frac{1}{2} \cdot 1}{q \cdot \frac{1}{2} \cdot 1 + (1-q) \cdot p(\theta = 1 \mid a_j = 0)} \\ &= \frac{q \cdot \frac{1}{2} \cdot 1}{q \cdot \frac{1}{2} \cdot 1 + (1-q) \cdot \frac{\frac{1}{2}(1-\alpha)}{\frac{1}{2}(1-\alpha) + \frac{1}{2}}} = \frac{q}{q + 2(1-q)\frac{1-\alpha}{2-\alpha}} \end{aligned}$$

The requirement that this expression exceeds $\frac{1}{2}$ boils down to the inequality in the statement of the proposition. The condition is not vacuous because there are values of $q < \frac{1}{2}$ that satisfy it, for any $\alpha > 0$. ■

Compare this result with an alternative specification in which $R_i = 1$ represents complete R-information. In this case, we would be unable to sustain the above equilibrium, because

$$p(a_j = 1 \mid \theta = R_i = 1) = p(R_j = 1 \mid R_i = 1) = q < \frac{1}{2}$$

Thus, the fact that $R_i = 1$ represents imperfect R-information enables an exaggerated belief that $a_j = 1$ conditional on $\theta = R_i = 1$, which facilitates good coordination. The reason is that player i's failure to account for the joint effects of θ and R_j on a_j leads him to regard θ and R_i as two conditionally independent signals of a_j, thus exaggerating the informative content of $\theta R_i = 1$. By itself, this is a correlation-neglect effect that appears in earlier work (e.g., in the context of voting behavior in Levy and Razin (2015)). The novelty here is that this effect is embedded in high-order reasoning: a particular realization of player i's archive-information leads him to neglect the correlation between player j's news-information.

4.3 A Hierarchical Type Space

In both the Prisoner's Dilemma example of the Introduction and the example of the previous sub-section, one 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' conventional signals and actions, and consequence variables. For each player i, define an m-tuple $(R_i^1, ..., R_i^m)$, $m \geq 2$, where R_i^1 is a collection of subsets of B; and for every $k = 2, ..., m$, R_i^k is a collection of subsets of $B \cup \{l_{R_i^1}, ..., l_{R_i^{k-1}}\} \cup \{l_{R_j^1}, ..., l_{R_j^{k-1}}\}$, where each of these subsets includes $l_{R_i^{k-1}}$ or $l_{R_j^{k-1}}$. Define $R_i = \cup_{k=1,...,m} R_i^k$.

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 "global game". For simplicity, assume that players have common R-information, which is distributed independently of θ. Assume in addition that players have complete N-information - i.e., $t_i = (\theta, R)$ for both $i = 1, 2$. The basic variables are θ, a_1, a_2. Suppose that for every $k = 1, 2, ...,$ R^k takes two values, 0 and 1, which are defined 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	$\varnothing$	$\{\{l_\theta, l_{a_1}, l_{a_2}, l_{R^1}\}\}$
3	$\varnothing$	$\{\{l_\theta, l_{a_1}, l_{a_2}, l_{R^1}, l_{R^2}\}\}$
$\vdots$	$\vdots$	$\vdots$
m	$\varnothing$	$\{\{l_\theta, l_{a_1}, l_{a_2}, l_{R^1}, ..., l_{R^{m-1}}\}\}$
$\vdots$	$\vdots$	$\vdots$

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 $p(R = k) = \delta(1 - \delta)^k$ 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 6 *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 of the 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

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

■

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 A Market Application

In this section I present a simple model of a market for a financial asset, where two large traders interact in the presence of exogenous external demand. As in conventional models of financial trade, traders may be diversely informed about the realized value of the traded asset. However, unlike standard models, the traders will also vary in their archive-information. I will refer to traders with better archive-information as sophisticated, and those with poorer archive-information as naive. The difference between the two types is that sophisticated traders' data encompasses the sophistication variable, whereas naive traders' data is limited to the variables that would suffice if all traders had rational expectations. The question is how this augmented model of diversely informed traders affects trading patterns.

Let us present the formal model. A single financial asset is traded in a market that is populated by two large traders, denoted 1 and 2. Each trader i simultaneously choose a position $a_i \in \mathbb{R}$ - i.e., short-selling is frictionless.

Exogenous external demand is given by the function $a_0(\pi) = 1 - \pi$, where π is the asset's price. There is a fixed supply of the asset, normalized to 1. The asset's price π is given by the market clearing condition

$$a_0(\pi) + a_1 + a_2 = 1$$

such that the market clearing price is $\pi = a_1 + a_2$ (i.e., the model treats negative and positive prices symmetrically). Trader i's payoff is $a_i(\theta - \pi)$, where $\theta \sim U\{-1, 1\}$ is the asset's intrinsic value. Plugging the external demand and market clearing condition, the trader i's payoff is reduced to

$$u_i(a_i, a_j) = a_i(\theta - a_i - a_j)$$

I now introduce the type space, which illustrates two features of the formalism: potential correlation between a player's archive-information and the conventional Harsanyi signal, and the possibility that one player's archive-information encompasses another player's archive-information. Let $s_i \in \{-1, 1\}$ denote a signal that trader i receives about the asset's intrinsic value. Let $R_i \in \{0, 1\}$ indicate trader i's R-information. For each trader i, $R_i = 1$ with independent probability λ. Trader i's N-information is the pair (R_i, s_i). The realization $R_i = 0$ corresponds to

$$\{\{\delta_\theta, \delta_{s_i}\}, \{\delta_\theta, \delta_{s_j}\}, \{\delta_{s_i}, \delta_{a_i}\}, \{\delta_{s_j}, \delta_{a_j}\}\}$$

and represents a naive trader. The realization $R_i = 1$ corresponds to

$$\{\{\delta_\theta, \delta_{s_i}, \delta_{\psi_i}\}, \{\delta_\theta, \delta_{s_j}, \delta_{\psi_j}\}, \{\delta_{s_i}, \delta_{a_i}, \delta_{\psi_i}\}, \{\delta_{s_j}, \delta_{a_j}, \delta_{\psi_j}\}\}$$

and represents a sophisticated trader.

Note that R_i treats the traders' variables symmetrically. The interpretation is that the game is anonymous, where traders are independently drawn from a large population of traders. The archival information a trader receives consists of anonymized data. For example, the realization $R_i = 0$ means that he receives data about the joint distribution of the asset's intrinsic value and the signal that any individual trader receives about it, as well as data about the joint distribution of individual traders' signal and actions.

I assume perfect positive correlation between the quality of traders' signals and their degree of sophistication. Specifically, for each trader i and every $\theta \in \{-1, 1\}$,

$$s_i = \begin{cases} \theta & \text{with probability} \quad \frac{1}{2}(1 + R_i) \\ -\theta & \text{with probability} \quad \frac{1}{2}(1 - R_i) \end{cases}$$

That is, a sophisticated trader receives a perfectly informative signal, whereas a naive trader receives a perfectly uninformative signal. The reason I focus on this specification is that by assuming perfect correlation between the two components of a trader's type, I ensure that there are as many trader types as in a rational-expectations benchmark (where all traders have perfect archive-information) - only the meaning of types changes. This makes the comparison with this benchmark as clean as possible.

The two possible realizations of traders' R-information satisfy the running intersection property. Therefore, $p_{R_i=0}$ factorizes p according to the perfect DAG

$$a_i \leftarrow s_i \leftarrow \theta \rightarrow s_j \rightarrow a_j$$

and $p_{R_i=1}$ factorizes p according to the perfect DAG

$$\begin{array}{ccccccccc} & & R_i & & & & R_j & & \\ & \swarrow & \downarrow & \nwarrow & & \nearrow & \downarrow & \searrow & \\ a_i & \leftarrow & s_i & \leftarrow & \theta & \rightarrow & s_j & \rightarrow & a_j \end{array}$$

This facilitates the derivation of trader i's subjective belief as a function of his type:

$$\begin{aligned} p_{R_i=0}(\theta, a_j \quad | \quad & R_i = 0, s_i) = p(\theta \mid s_i) \sum_{s_j} p(s_j \mid \theta) p(a_j \mid s_j) \\ = \quad & \sum_{R_i} p(R_i \mid s_i) p(\theta \mid R_i, s_i) \sum_{s_j} \sum_{R_j} p(R_j) p(s_j \mid \theta, R_j) \sum_{R'_j} p(R'_j \mid s_j) p(a_j \mid s_j, R'_j) \end{aligned}$$

and

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

In a standard model, $p_{R_i=0}$ would coincide with rational expectations, because the conditional-independence properties $s_i \perp s_j \mid \theta$, $a_i \perp \theta \mid s_i$ and $a_j \perp \theta \mid s_j$ would be satisfied by assumption. However, in the present model, R_i is a confounding variable that affects both the distribution of s_i conditional on θ and (potentially) the distribution of a_i conditional on s_i. Therefore, the above conditional-independence properties need not hold, which means that $p_{R_i=0}$ can depart from rational expectations. In contrast, $p_{R_i=1}$ coincides with rational expectations.

Let us now turn to equilibrium characterization. The equilibrium is unique, symmetric and in pure strategies; proving this is conventional, and because it would make the proof notationally messier, I take it for granted. Denote by x_s^R the equilibrium action taken by a trader whose archive-information and news-information are R and (R, s), respectively.

Proposition 7 *Equilibrium actions are given by*

$$\begin{aligned} x_s^1 &= \frac{s}{2+\lambda} \\ x_s^0 &= \frac{s}{2+\lambda} \cdot \frac{2\lambda + \lambda^2 - \lambda^3}{2 + \lambda^2 - \lambda^3} \end{aligned}$$

Proof. Trader i's subjective expected utility from a_i given his type (R_i, s_i)

is

$$a_i \cdot [E_{R_i}(\theta \mid s_i, R_i) - a_i - E_{R_i}(a_j \mid s_i, R_i)]$$

The standard first-order condition for the subjective optimality of $x_{s_i}^{R_i}$ is

$$x_{s_i}^{R_i} = \frac{1}{2} [E_{R_i}(\theta \mid s_i, R_i) - E_{R_i}(a_j \mid s_i, R_i)] \tag{4}$$

Let us first obtain expressions for $E_{R_i=1}(\theta \mid s_i, R_i = 1)$ and $E_{R_i=1}(a_j \mid s_i, R_i = 1)$. Because $R_i = 1$ induces rational expectations, these expressions coincide with $E(\theta \mid s_i, R_i = 1)$ and $E(a_j \mid s_i, R_i = 1)$, which are given as follows. Because a trader i with $R_i = 1$ receives a perfectly informative signal,

$$E(\theta \mid s_i, R_i = 1) = s_i$$

In addition, this trader predicts correctly trader j's action as a function of his type. Therefore,

$$E(a_j \mid s_i, R_i = 1) = \lambda \cdot x_{s_i}^1 + (1 - \lambda) \cdot (\frac{1}{2} x_{s_i}^0 + \frac{1}{2} x_{-s_i}^0)$$

Here we make use of the assumption that if trader j is sophisticated (i.e. $R_j = 1$), his signal is perfectly informative, such that $s_j = \theta = s_i$ with probability one. Plugging this into (4), we obtain

$$x_{s_i}^{R_i} = \frac{1}{2} \left[s_i - \lambda \cdot x_{s_i}^1 - (1 - \lambda) \cdot (\frac{1}{2} x_{s_i}^0 + \frac{1}{2} x_{-s_i}^0) \right]$$

Let us now derive expressions for $E_{R_i=0}(\theta \mid s_i, R_i = 0)$ and $E_{R_i=0}(a_j \mid s_i, R_i = 0)$. I will do so for $s_i = 1$ - the case of $s_i = -1$ is symmetric. First,

$$p(\theta = 1 \mid s_i = 1) = \frac{\frac{1}{2}[\lambda + (1 - \lambda)\frac{1}{2}]}{\frac{1}{2}} = \frac{1 + \lambda}{2}$$

such that

$$E_{R_i=0}(\theta \mid s_i = 1) = \frac{1 + \lambda}{2} \cdot 1 + \frac{1 - \lambda}{2} \cdot (-1) = \lambda$$

To calculate $E_{R_i=0}(a_j \mid s_i, R_i = 0)$, we first need to calculate $E_{R_i=0}(a_j \mid \theta)$ for both $\theta = 1, -1$. I will do so for $\theta = 1$ - the case of $\theta = -1$ is symmetric:

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

Let us derive the terms in this expression. First,

$$\begin{aligned} p(s_j &= 1 \mid \theta = 1) = \lambda + (1-\lambda)\frac{1}{2} = \frac{1+\lambda}{2} \\ E(a_j &\mid s_j) = p(R_j = 1 \mid s_j) x^1_{s_j} + p(R_j = 0 \mid s_j) x^0_{s_j} \end{aligned}$$

Now, it is easy to verify that because the ex-ante distribution over s_j as well as the distribution over s_j conditional on R_j are both symmetric, $p(R_j = 1 \mid s_j) = p(R_j = 1) = \lambda$ for all s_j. Therefore,

$$E_{R_i=0}(a_j \mid \theta = 1) = \frac{1+\lambda}{2} \cdot [\lambda x^1_1 + (1-\lambda) x^0_1] + \frac{1-\lambda}{2} \cdot [\lambda x^1_{-1} + (1-\lambda) x^0_{-1}]$$

In the same manner, we can derive

$$E_{R_i=0}(a_j \mid \theta = -1) = \frac{1+\lambda}{2} \cdot [\lambda x^1_{-1} + (1-\lambda) x^0_{-1}] + \frac{1-\lambda}{2} \cdot [\lambda x^1_1 + (1-\lambda) x^0_1]$$

Now,

$$E_{R_i=0}(a_j \mid s_i, R_i = 0) = p(\theta = 1 \mid s_i) E_{R_i=0}(a_j \mid \theta = 1) + p(\theta = -1 \mid s_i) E_{R_i=0}(a_j \mid \theta = -1)$$

Plugging the relevant expressions in (4), we obtain

$$x^0_1 = \frac{1}{2}\left[\lambda - \left(\frac{\lambda}{2}(1+\lambda^2)x^1_1 + \frac{\lambda}{2}(1-\lambda^2)x^1_{-1} + \frac{1-\lambda}{2}(1+\lambda^2)x^0_1 + \frac{1-\lambda}{2}(1-\lambda^2)x^0_{-1}\right)\right]$$

and a similar expression for x^0_{-1}. It is then immediate that $x^0_{-1} = x^0_1$ and $x^1_{-1} = x^1_1$, which gives the solution. ■

It is instructive to compare this result to the rational-expectations benchmark in which all traders have complete archive-information, a fraction λ of traders are perfectly informed of θ and a fraction $1-\lambda$ of traders are perfectly uninformed. In that alternative model, perfectly uninformed traders would be completely inactive, and a perfectly informed trader who receives the signal s would play $s/(2+\lambda)$. Thus, the equilibrium behavior of sophisticates in our model is exactly the same as in the rational-expectations benchmark. The reason is that because naive traders' signals are independent of θ and because $x_1^0 = x_{-1}^0$, their expected position conditional on the sophisticated trader's signal is zero. As far as the sophisticated trader is concerned, naive traders might as well be completely inactive.

The difference between the two models lies in the naive traders' behavior, as they choose non-zero quantities. The reason naive traders are active is that they condition their belief over θ on their signal, failing to account for its accuracy. In other words, they respond to the signal as if it is partially informative. Furthermore, a naive trader neglects the role of his opponent's archive-information on his behavior, and therefore he ends up underestimating the opponent's response to θ. This is a force that impels the naive trader to respond even more aggressively to his signal, and choose a higher absolute quantity that he would have chosen had he understood the mapping from θ to the opponent's expected position. Thus, overall trading activity in the market - measured by the ex-ante expected absolute quantity taken by each trader - is higher than in the rational-expectations benchmark.

Naive traders' equilibrium quantity is lower in absolute terms than sophisticated traders. This is a net result of two opposing forces. On one hand, naive traders' belief over θ is less responsive to their signal, because they regard it as partially informative (in reality it is entirely uninformative), whereas sophisticated traders (rightly) regard their own signal as perfectly informative. This is a force that curbs the naive trader's activity. On the other hand, as already mentioned, because a naive trader neglects the role of his opponent's archive-information on his behavior, he ends up underestimating the opponent's response to θ, which is a force that impels the naive trader to be more aggressive and partially offsets the first force.

As the fraction of sophisticated traders λ increases, their position increases in absolute terms, while naive traders' position decreases in absolute terms. The ex-ante expected absolute quantity that a trader chooses in equilibrium is

$$\frac{1}{2+\lambda} \cdot \left[\lambda \cdot 1 + (1-\lambda) \cdot \frac{2\lambda + \lambda^2 - \lambda^3}{2 + \lambda^2 - \lambda^3} \right] = \frac{4\lambda - \lambda^2 - \lambda^3}{(2+\lambda)(2+\lambda^2-\lambda^3)}$$

This expression is hump-shaped. It attains the value 0 at $\lambda = 0$ and the value $\frac{1}{3}$ at $\lambda = 1$. It reaches a maximum of approximately 0.344 at $\lambda^* \approx 0.81$. Thus, the fraction of naive traders in the population has a slightly non-monotone effect on overall activity.

6 On the Formalism's Expressive Power

The archive-information formalism enables us to describe players whose subjective beliefs distort objective steady-state distributions, as a result of parsimonious extrapolation from partial knowledge of steady-state correlations. While the numerous examples in this paper illustrated the formalism's expressive power, this power is not unlimited - i.e., certain belief distortions lie outside its scope. E.g., consider a steady-state distribution in which all variables are statistically independent. Then, whatever the player's R-information, his subjective belief p_R must coincide with the objective distribution: the maximal-entropy extrapolation rule forbids the player from dreaming up correlations that do not actually exist.

In this section I carry out a systematic exploration of the expressive scope of the archive-information formalism, in a simple class of environments. Let the relevant variables be the state of nature θ; the players' actions a_1, a_2; and their archive-information R_1, R_2. Thus, each player's archive-information is defined by a single variable. The payoff-relevant variables are θ, a_1, a_2. Player i's N-information is $t_i = (\theta, R_i)$. That is, players have complete news-information of the state of Nature and their own archive-information (the latter part of this assumption is redundant, given the restrictions I am about to impose).

Let us focus on steady-state distributions p over $\theta, a_1, a_2, R_1, R_2$ that satisfy $a_1, R_1 \perp a_2, R_2 \mid \theta$. In addition, impose the following restrictions on the possible realizations of R_i in p. First, $\{\delta_\theta, \delta_{a_j}\} \subseteq N(R_i) \subseteq \{\delta_\theta, \delta_{a_j}, \delta_{R_j}\}$. That is, player i's archive-information only covers the state of Nature and the variables that describe his opponent. Second, R_i satisfies the running intersection property, such that p_{R_i} has a perfect-DAG representation (see Section 3). The first restriction implies that we can write player i's belief regarding a_j conditional on his type as $p_{R_i}(a_j \mid \theta)$. The second restriction implies that the player's subjective belief does not distort the steady-state marginal distributions over individual variables. In particular, $p_{R_i}(a_j) \equiv p(a_j)$ and $p_{R_i}(\theta) \equiv p(\theta)$.

To gauge the expressive scope of this formalism, we need to compare it to a benchmark model that allows for less structured belief distortion. Denote by Θ and A_i the set of states of Nature and let A denote the set of actions available to each player. Let C be the set of all functions $c : \Delta(\Theta \times A) \rightarrow \Delta(\Theta \times A)$ that preserve marginals over θ and a. The interpretation is that $c_i \in C$ is player i's "cognitive type" - it defines a systematic belief distortion function, which maps any objective distribution $\tilde{p}(\theta, a_j)$ into a subjective belief, denoted $\tilde{p}_{c_i}(\theta, a_j)$. A steady-state distribution $\tilde{p}$ in this model is thus defined over the variables $\theta, a_1, a_2, c_1, c_2$, such that given the realization θ, c_i, player i's belief regarding player j's action is $\tilde{p}_{c_i}(a_j \mid \theta)$. I will refer to such a distribution $\tilde{p}$ as an "unstructured steady-state distribution" (USSD).

The class of USSDs contains the set of steady-state distributions that can arise in the above archive-information model. The reason is that every R_i in the permissible domain above induces a (unique) cognitive type, which maps any joint distribution $p(\theta, a_j)$ into $p_{R_i}(\theta, a_j)$ while preserving the marginals over the individual variables θ and a_j. Let $c(R_i)$ denote the cognitive type induced by R_i. The set C is a continuum, whereas there is only a finite number of possible realizations of R_i. Therefore, the set of cognitive types c for which there exists R_i such that $c(R_i) = c$ has measure zero.

It follows that any steady-state distribution p over $\theta, a_1, a_2, R_1, R_2$ in the archive-information model can be mapped into a USSD $\tilde{p}$:

$$\tilde{p}(\theta, c_i, a_i) \equiv \sum_{R_i \mid c(R_i) = c_i} p(\theta, R_i, a_i)$$

We will say that the mapping is *tight* if $c(R_i) \neq c(R'_i)$ for every distinct R_i, R'_i in the support of p - i.e., distinct realizations of player i's archive information correspond to distinct cognitive types in the USSD.

I now present two simple results that characterize the steady-state distributions in the unstructured model that can be induced in this manner by steady-state distributions in the archive-information model described above. The results will make use of two special cognitive types. Let c^{RE} be the cognitive type that never distorts $p(a_j \mid \theta)$ - i.e., $p_{c^{RE}}(a_j \mid \theta) \equiv p(a_j \mid \theta)$. This type is indistinguishable from a player with rational expectations. Let c^{FC} be a "fully cursed" cognitive type that completely fails to account for the correlation between θ and a_j - i.e., $p_{c^{FC}}(a_j \mid \theta) \equiv p(a_j)$. The first result in this section establishes that these are the only cognitive types that can appear in a USSD that is induced by an archive-information steady-state distribution in which players' archive-information is *independent of the state of Nature.*

Proposition 8 *Let p be a steady-state distribution in the archive-information model in which $R_i \perp \theta$ for every i. Then, the only cognitive types in the support of the induced USSD $\tilde{p}$ are c^{RE} and c^{FC}.*

Proof. Consider some R_i with $p(R_i) > 0$. Suppose that $\{\delta_\theta, \delta_{a_j}\}$ is contained in some set in R_i. Then, $\{\delta_\theta, \delta_{a_j}\}$ is a clique in a perfect DAG that represents R_i. Therefore, $p_{R_i}(\theta, a_j) \equiv p(\theta, a_j)$, such that $p_{R_i}(\theta \mid a_j) \equiv p(\theta \mid a_j)$ - i.e., $c(R_i) = c^{RE}$. Now suppose that no set in R_i contains $\{\delta_\theta, \delta_{a_j}\}$. Then, $\{\delta_\theta, \delta_{a_j}\}$ is not a clique in any perfect DAG that represents R_i. This means that p_{R_i} accounts for the correlation between θ and a_j only in-so-far as these variables are correlated with R_j. However, by assumption, $R_j \perp \theta$. Therefore, $p_{R_i}(\theta, a_j) \equiv p(\theta)p(a_j)$, such that $p_{R_i}(a_j \mid \theta) \equiv p(a_j)$ - hence, $c(R_i) = c^{FC}$. ■

Thus, a USSD in which independent cognitive types has an archive-information foundation only when the only cognitive types it admits are the rational-expectations and fully-cursed types. This means that in order to enrich the expressive scope of the archive-information model, we must allow players' archive-information to be correlated with the state of Nature.

The following result allows for correlation between players' archive-information and the state of Nature.

Proposition 9 *Let $\tilde{p}$ be a USSD that is tightly induced by a steady-state distribution p in the archive-information model. Then, for each player i, $\tilde{p}$ assigns positive probability to at most three cognitive types: c^{RE}, c^{FC}, and a type c_i^* defined by*

$$\tilde{p}_{c_i^*}(a_j \mid \theta) = \sum_{c_j \in \{c^{RE}, c^{FC}, c_j^*\}} p(c_j \mid \theta) p(a_j \mid c_j)$$

Proof. Consider a steady-state distribution p in the archive-information model. The restrictions we have imposed on players' archive-information imply that R_i can get at most five values:

$$\begin{aligned}
R^1 &= \{\{\delta_\theta\}, \{\delta_{a_j}\}\} \\
R^2 &= \{\{\delta_\theta, \delta_{R_j}\}, \{\delta_{a_j}\}\} \\
R^3 &= \{\{\delta_\theta\}, \{\delta_{R_j}, \delta_{a_j}\}\} \\
R^4 &= \{\{\delta_\theta, \delta_{R_j}\}, \{\delta_{R_j}, \delta_{a_j}\}\} \\
R^5 &= \{\{\delta_\theta, \delta_{R_j}, \delta_{a_j}\}\}
\end{aligned}$$

The realization R^5 induces rational expectations - i.e., $p_{R^{(5)}}(a_j \mid \theta) \equiv p(a_j \mid \theta)$. Therefore, $c(R^5) = c^{RE}$. The realizations R^1, R^2, R^3 all induce "fully cursed" beliefs - i.e. $p_R(a_j \mid \theta) \equiv p(a_j)$. That is, $c(R^1) = c(R^2) = c(R^3) = c^{FC}$. Therefore, if a USSD $\tilde{p}$ is *tightly* induced by p, it must be the case that p assigns positive probability to at most one of these three realizations of R_i. The realization R^4 induces the belief $p_{R^4}(a_j \mid \theta) = \sum_{R_j} p(R_j \mid \theta)p(a_j \mid R_j)$. We have established that p assigns positive probability to at most three values of R_j, where two of these values have $c(R_j) = c^{FC}$ and $c(R_i) = c^{RE}$. It follows that $c(R_i) = c^*$ for this value of R_i. ■

Thus, allowing for correlation between players' archive-information and the state of Nature (while demanding a *tight* mapping between the two models) slightly enlarges the set of USSDs that can be generated by our formalism. Each player i has at most one additional cognitive type. Moreover, the particular belief distortion administered by this additional type is entirely pinned down by the joint distribution over θ, c_j, a_j. Therefore, the archive-information model imposes a strong link between the kinds of belief distortion that can occur in a steady-state distribution and how they correlate with other variables.

The results in this section paper demonstrate that the archive-information formalism imposes a great deal of structure on the patterns of subjective belief distortions that can be displayed in a steady-state distribution. This can be viewed as a limitation of the formalism because it highlights its restricted expressive scope. Alternatively, it can be hailed as a merit, because very often modeling frameworks (especially in the field of bounded rationality are criticized for being able to "explain everything". From this perspective, the restrictions on the patterns of belief distortion that can arise in this framework is a feature rather than a bug. Extending this characterization to the case in which players' archive-information is defined by multiple variables (as in the hierarchical construction of Section 4.3) is left for future work.

7 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 their description of their *type*. Importantly, it 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. I plan to explore this direction in follow-up work.

References

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

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

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

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

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

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

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

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

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

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

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

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

[13] Koller, D. and N. Friedman (2009), Probabilistic Graphical Models: Principles and Techniques. MIT press.

[14] 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.

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

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

[17] Pearl, J. (2009), Causality: Models, Reasoning and Inference, Cambridge University Press, Cambridge.

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

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

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

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

[22] Spiegler, R. (2017b), Can Agents with Causal Misperceptions be Systematically Fooled? CEPR Discussion Paper 11379.