Modeling Players with Random "Data Access"*

Ran Spiegler[†]

October 12, 2021

Abstract

I present an approach to static equilibrium modeling with non-rational expectations, which is based on enriching players' typology. A player is characterized by his "data access", consisting of: (*i*) "news access", which corresponds to a conventional signal in the Harsanyi model, and (*ii*) "archival access", a novel component representing the player's piecemeal knowledge of steady-state correlations. Drawing on prior literature on correlation neglect and coarse reasoning, I assume the player extrapolates a well-specified probabilistic belief from his "archival data" according to the maximum-entropy criterion. I show with a series of examples how this formalism extends our ability to represent and analyze strategic interactions without rational expectations.

---

*Earlier versions of the paper were circulated under the title "News and Archival Information in Games". Financial support by ERC Advanced Investigator grant no. 692995 is gratefully acknowledged. I thank Kfir Eliaz, Nathan Hancart, Ariel Rubinstein, Heidi Thysen, the editor of this journal and two referees, as well as seminar audiences at Tel Aviv, Hebrew and New York Universities, for helpful comments. I also thank Briq for its hospitality during the writing of this paper.

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

1 Introduction

Conventional game theory distinguishes between two kinds of knowledge that players in a static game may have. The first kind concerns the *realization* of exogenous variables (e.g., players' preferences). The second kind concerns the joint statistical behavior of exogenous and endogenous variables (e.g., how players' bids in an auction vary with their preferences). To use a journalistic metaphor, knowledge of the first kind is akin to a newsflash about an economic indicator, whereas the second kind is like the data that a reporter collects when he digs the newspaper's archives for evidence about what this indicator signified in past episodes. Accordingly, I refer to the two kinds of knowledge as *news information* and *archival information*.

The model of Bayesian games, as conventionally practiced by economists, offers a rich description of players' news information, while leaving the description of players' archival information to the solution concept. "One-shot" solution concepts like rationalizability implicitly assume that players lack archival information about endogenous variables. At the other extreme, Nash equilibrium presumes that players have complete archival information and therefore know the steady-state joint distribution of all variables; in other words, they have "rational expectations."

The literature on equilibrium behavior with non-rational expectations (Osborne and Rubinstein (1998), Jehiel (2005), Jehiel and Koessler (2008), Esponda (2008), Esponda and Pouzo (2016)) has proposed solution concepts that retain the steady-state approach of Nash equilibrium, while replacing complete archival information with a notion of limited feedback about the steady-state distribution and a notion of how players extrapolate a belief from their feedback. I provide a detailed literature review in Section 6. At this stage, it suffices to say that virtually all prior proposals assume that players' feedback limitations are fixed. And neither provides a model of players' imperfect knowledge (of either kind) regarding their opponents' archival information.

Yet, it is easy to think of real-life situations that call for such a model. A market agent's "type" may be defined in terms of both kinds of knowledge, which can be correlated - an agent with better access to reliable news sources is also likely to have better access to archival information. We can also talk meaningfully about one player's news information regarding another player's archival information. To use a military-intelligence example, suppose Army 1 hears (from a dubious source) that Army 2 has just gained access to archival records of Army 1's force deployment in various weather conditions. Alternatively, one player may have partial archival information about another player's archival access - e.g., with some probability army 1 obtains a computer file documenting Army 2's archival access in various situations.

In this paper I present a new typology of non-rational expectations in static strategic interactions, which combines news and archival information. For expositional simplicity, I restrict attention to two-player interactions. A state of the world is described by the realization of a collection of *variables*. There is an objective distribution p over states, interpreted as a *steady-state distribution*. Importantly, the description of a state also includes the realization of *endogenous* variables (players' actions, final allocations), such that p has both exogenous and endogenous components.[1]

A player is characterized by his "*data access*", which has two components. The player's "*news access*" specifies the exogenous variables whose realization he learns before taking his action. The realization of these variables constitutes the player's "*news information*". This corresponds to the usual notion of a Harsanyi signal, formulated in a slightly unusual manner. The novelty lies in the second component, namely the player's "*archival access*", which is defined as a collection of subsets of variables. This means that the player gets to learn the marginal of p over each of these subsets of variables. Thus, rather than learning the entire joint distribution p (as in the case of Nash equilibrium), the player has piecemeal knowledge of p in the form of

[1] This idea is borrowed from Aumann (1987) but plays a very different role.

certain marginals over subsets of variables. These marginals constitute the player's "*archival information*". I assume that the player's payoff function is always measurable with respect to the variables in his archival data.

In this formulation of archival information, I have in mind situations in which the player lacks an understanding of the rules of the game. He interacts *once*, after getting statistical feedback about past interactions of other agents. The distribution p records the frequencies of realizations of all relevant variables in these past interactions. The player's understanding of long-run statistical regularities - and therefore his perception of the mapping from actions to payoffs - is based *purely* on the learning feedback given by his archival information. This is in the spirit of concepts like Berk-Nash or Analogy-Based Expectations Equilibrium (see Section 6).

The following table visualizes the notion of random archival access in strategic situations. Suppose player 1 can receive statistical data about the joint steady-state distribution of three variables. There are three datasets, each involving a different combination of variables. E.g., dataset I describes the joint distribution of the state of Nature and player 2's action (more precisely, the actions of agents who assumed the role of player 2 in past interactions). Dataset II describes the joint distribution of the state of Nature and player 2's archival access (more precisely, the access enjoyed by agents who assumed the role of player 2 in past interactions). Player 1's archival access is random: he may get access to different combinations of the three datasets. When he gets access to multiple datasets, his task is to extrapolate a belief from his different pieces of data; he has no other basis for forming his beliefs.

	State of Nature	Player 2's archival access	Player 2's action
Dataset I	✓		✓
Dataset II	✓	✓	
Dataset III		✓	✓

One of the variables on which player 1 can get data is the *archival access of player* 2, which is a random variable as well. This captures one of the "military intelligence" examples mentioned above: it gives player 1 archival data about how the archival access of agents who assume player 2's role varies with the state of Nature. Although this notion of "archival information about another player's archival access" connotes the kind of introspective high-order beliefs we sometimes encounter in the theory of Bayesian games, players in my framework do not use introspection to form beliefs: they rely entirely on their partial statistical data.

The player in my framework forms his beliefs in two stages. First, he extrapolates a subjective probabilistic belief from his archival information, thus potentially distorting p. There are many extrapolation rules one could employ. Yet, 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 lack evidence*. Such "correlation neglect" has been studied theoretically (e.g., Levy and Razin (2015)) and experimentally (e.g., Eyster and Weizsacker (2016), Enke and Zimmermann (2017).[2] To capture this motive, I assume that the player's belief is the distribution (over variables in his archive) that *maximizes Shannon entropy*, subject to being consistent with the marginals he learns. This subsumes several existing notions of non-rational expectations as special cases.

The second stage of the belief-formation procedure is conventional. The player conditions his first-stage belief on his news information via Bayes' rule - just as he would in the standard model - and thus forms a subjective mapping from actions to payoff-relevant outcomes. While superficially standard, it produces non-standard effects (such as play of dominated actions) when combined with the first stage. Equilibrium is also defined conventionally, in the familiar manner of trembling-hand perfection: If an action fails to

[2] Correlation neglect has been mentioned as a culprit in professional analysts' failure to predict political and economic events (see Hellwig (2008) and https://fivethirtyeight.com/features/the-real-story-of-2016).

maximize a player's subjective expected utility given his news information, it is played with vanishing probability. Trembles are interpreted as blind experimentation, and they ensure that the procedure's second stage does not involve conditioning on zero-probability events.

My task in this paper is to demonstrate how this framework expands our ability to represent and analyze economic situations populated by agents with random data access.

2 A Motivating Example: "Navigation Apps" in a Congestion Game

I present an example that gives a foretaste of the formalism, before describing it in full generality. Consider a congestion situation in which multiple car drivers simultaneously choose driving routes from a fixed starting point to a fixed destination. Each driver's objective is to minimize expected driving time. Because of negative congestion externalities, the driving time in a given route segment is weakly increasing in the number of drivers who take it.

The standard solution concept for describing drivers' behavior is Nash equilibrium. However, this concept presumes that drivers know the equilibrium mapping from their choice of route to the expected total driving time. Instead, let us assume that drivers vary in the quality of their driving-time predictions. Some drivers are equipped with a "*navigation app*" that bases its predictions on fine statistical data and therefore adjusts them to current conditions (which may include weather, timing, construction works, etc.). Other drivers lack a navigation app; their predictions are based on coarser statistical data. Moreover, the composition of the two types in the population of drivers can vary with external conditions. The data-access language enables us to formalize this situation: whether a driver is equipped with the app effectively defines his archival access.

The following congestion game is a simple instance of the well-known

Braess Paradox (Braess (1968)), taken from Osborne (2004, pp. 34-35)). The following directed graph provides a template for the game:

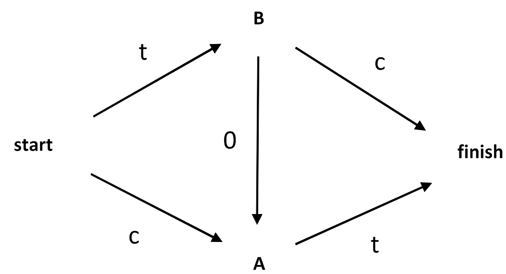

There are two drivers, denoted 1 and 2. They begin their drive simultaneously at the *start* node and must end it at the *finish* node. Each directed link represents a route segment, and the number written near it indicates the travel time in this segment, where t is the number of drivers who take it and $c \in (2, 2.5)$ is a constant. Each driver has three strategies: (A) go to A and from there to *finish*; (B) go to B and from there directly to *finish*; (BA) go to B, then to A and then to *finish*. The total travel time for driver i is denoted z_i. The realization of (z_1, z_2) is a function of the players' action profile (a_1, a_2), as given by the following bi-matrix:

$a_1 \backslash a_2$	A	B	BA
A	$c+2, c+2$	$c+1, c+1$	$c+2, 3$
B	$c+1, c+1$	$c+2, c+2$	$c+2, 3$
BA	$3, c+2$	$3, c+2$	$4, 4$

(1)

Driver i's payoff is $-z_i$.

The action BA is strictly dominant. However, (BA, BA) is Pareto-dominated by the pairs (A, B) and (B, A). Furthermore, if the segment

$B \to A$ were *blocked* and the strategy BA thus became infeasible, (A, B) and (B, A) would be the Nash equilibria in the reduced 2×2 game. This is the Braess Paradox.

I now enrich this interaction with random archival access. I use p to denote an objective steady-state distribution over all relevant variables (to be introduced as I go along). Let $\theta \in \{0, 1\}$ be a state of Nature. The two states are equally likely. Players' archival access is a deterministic function of θ. When $\theta = 1$, each player i receives archival data that enables him to learn the objective joint distribution $(p(\theta, a_i, z_i))$. When $\theta = 0$, each player i receives archival data that only enables him to learn the objective joint distribution $(p(a_i, z_i))$. Note that I use the same symbol p to describe any marginal distribution that is induced by the fully defined steady-state distribution. This abuse of notation will serve me throughout the paper.

I use the notations $r_i = \{\{\boldsymbol{\theta}, \mathbf{a}_i, \mathbf{z}_i\}\}$ and $r_i = \{\{\mathbf{a}_i, \mathbf{z}_i\}\}$ to describe these realizations of player i's archival access. The boldface notation enables me to distinguish between a variable and its label: the notation $\mathbf{x}$ represents the label of the variable x.

The interpretation is as follows. When $\theta = 1$, players have a navigation app that provides fine archival data and thus enables them to predict the total driving time in each route as a function of θ. When $\theta = 0$, players lack access to the app and therefore have coarse archival data, which only enables them to learn the average driving time in each route, aggregating over θ. This captures one of the modeling framework's features described in the Introduction - namely, the possibility that players lack archival data about a relevant variable. Since the variable in question determines players' archival access, the example illustrates another feature: the possibility that players have partial archival access to variables that define other players' archival access. Section 4 will examine more elaborate instances (they will also involve the aspect of maximum-entropy extrapolation, which is missing from the current example).

As to players' news access, assume it is constant: each driver is always informed of the realized state of Nature and his own characteristics. The list of relevant variables is thus $\theta, r_1, r_2, a_1, a_2, z_1, z_2$, and p is interpreted as a steady-state joint distribution over these variables.

Turning to players' belief formation, the key assumption here is that players only form beliefs over variables on which they have archival data. Thus, when $r_i = \{\{\boldsymbol{\theta}, \mathbf{a}_i, \mathbf{z}_i\}\}$, driver i's first-stage, unconditional belief is $(p(\theta, a_i, z_i))$; and his second-stage, conditional belief over payoff-relevant outcomes given his news information and his action is $(p(z_i \mid \theta, a_i))$. When $r_i = \{\{\mathbf{a}_i, \mathbf{z}_i\}\}$, driver i's first-stage belief is $(p(a_i, z_i))$ and his second-stage belief is $(p(z_i \mid a_i))$. Since the driver lacks archival access to θ, he is unable to condition on it.[3]

Player i's strategy is given by the conditional distribution $(p(a_i \mid \theta))$. The strategy profile constitutes an ε-equilibrium if, whenever $p(a_i \mid \theta) > \varepsilon$, a_i maximizes player i's expected payoff with respect to his subjective belief given θ. An equilibrium is the limit of a sequence of ε-equilibria with $\varepsilon \to 0$.

Claim 1 *There exist equilibria that depart from the Nash-equilibrium benchmark.*

In fact, there are exactly two such equilibria. Both are symmetric, and in each of them, the drivers' strategy takes the following form: *choose BA when $\theta = 1$, and play A and B with some probability α each when $\theta = 0$.* Since this is not a *fully* mixed strategy, we need to allow for trembles. Accordingly, suppose each driver follows this strategy with near certainty, and mixes uniformly over all three actions with the vanishing remaining probability. For brevity, I derive one of these two equilibria in detail and merely state the other.

[3] Since the driver never has archival access to r_i in this example, he is unable to condition on it. However, it can be shown that this distinction is moot in this example.

Consider the realization $\theta = 1$. In this case, driver i's subjective conditional belief over z_i is $(p(z_i \mid \theta, a_i))$. Since this is the correct distribution over z_i given θ, the driver will choose the dominant action BA, which is consistent with our guess. Now consider the realization $\theta = 0$. In this case, driver i's subjective conditional belief over z_i is $(p(z_i \mid a_i))$. Therefore, his subjective conditional expectation of z_i is

$$E(z_i \mid a_i) = \sum_{\theta} p(\theta \mid a_i) E(z_i \mid \theta, a_i) = \sum_{\theta} p(\theta \mid a_i) \left[\sum_{a_j} p(a_j \mid \theta) \cdot z_i(a_i, a_j) \right]$$

where $z_i(a_1, a_2)$ is given by (1).

Let us calculate each term in this expression, focusing on $i = 1$ without loss of generality. First,

$$\begin{aligned} E(z_1 &\mid \theta = 1, a_1 = A) = E(z_1 \mid \theta = 1, a_1 = B) \approx c + 2 \\ E(z_1 &\mid \theta = 1, a_1 = BA) \approx 4 \end{aligned}$$

because by assumption, player 2 chooses BA with near certainty when $\theta = 1$. The $\approx$ sign is due to the trembles.

Second,

$$\begin{aligned} E(z_1 &\mid \theta = 0, a_1 = A) = E(z_1 \mid \theta = 0, a_1 = B) \approx \alpha \cdot (c + 1) + (1 - \alpha) \cdot (c + 2) \\ E(z_1 &\mid \theta = 0, a_1 = BA) \approx 2\alpha \cdot 3 + (1 - 2\alpha) \cdot 4 \end{aligned}$$

because by assumption, player 2 approximately plays A and B with probability α each and BA with probability $1 - 2\alpha$ when $\theta = 0$.

Finally, given the drivers' presumed strategy and the prior distribution over θ,

$$\begin{aligned} p(\theta &= 0 \mid a_1 = A) = p(\theta = 0 \mid a_1 = B) \approx 1 \\ p(\theta &= 0 \mid a_1 = BA) \approx \frac{\frac{1}{2} \cdot (1 - 2\alpha)}{\frac{1}{2} \cdot (1 - 2\alpha) + \frac{1}{2} \cdot 1} \end{aligned}$$

Plugging these expressions into the formula for $E(z_i \mid a_i)$, we obtain

$$\begin{aligned} E(z_1 &\mid a_1 = A) = E(z_1 \mid a_1 = B) \approx c + 2 - \alpha \\ E(z_1 &\mid a_1 = BA) \approx 4 - \frac{\alpha(1 - 2\alpha)}{1 - \alpha} \end{aligned}$$

It follows that for every $c \in (2, 2.5)$, we can sustain an equilibrium in which $\alpha = \frac{1}{2}$ - namely, drivers never take the action BA when $\theta = 0$. This claim holds because $c < 2.5$ implies

$$c + 2 - \frac{1}{2} < 4 - \frac{\frac{1}{2} \cdot (1 - 2 \cdot \frac{1}{2})}{1 - \frac{1}{2}}$$

Note that when $\theta = 0$, driver's objective expected driving time is $c + 1.5 < 4$ - i.e., *the Braess paradox is alleviated.*

The intuition for this equilibrium is that when driver 1 lacks a navigation app (and therefore has limited archival access), he measures the correct correlation between the route he takes and total driving time, but he fails to condition this correlation on the state of Nature (because he lacks archival data about it). His calculation is tantamount to falsely interpreting the correlation between a_1 and z_1 as a pure *causal* effect of the former on the latter. In reality, this correlation is partly due to confounding by θ. When $\theta = 0$, the driver effectively ignores this confounder. As a result, he may end up believing that if he switches from BA to A or B, this will have a beneficial effect on his driving time. Thus, the underlying reason for the false-causation effect is that the player sometimes lacks archival data about a variable that confounds the relation between his action and another payoff-relevant variable.

Another noteworthy feature of this example is that the variation in players' actions is entirely due to the fluctuations in their data access. Because these fluctuations are correlated, the actions end up being correlated as well, yet one player type misconstrues this correlation. Of course, correlated actions could arise from other sources, e.g., correlated payoff shocks. What is special about the present example is that there is no payoff uncertainty; the source of correlated actions is the correlated randomness of players' knowledge of correlations.[4]

In the second symmetric equilibrium of the above-described form, drivers mix over *all* three actions when $\theta = 0$, such that $E(z_1 \mid a_1)$ is the same for all a_1. The solution to this indifference condition is

$$\alpha = 1 - \frac{c}{2} + \frac{1}{2}\sqrt{c^2 - 4}$$

In addition, the Nash-equilibrium outcome can also be sustained in equilibrium, with an appropriate choice of trembles. It can be shown that no other equilibria exist. Thus, while there is a unique Nash equilibrium in the original game, its archival-data extension gives rise to multiple equilibria.

3 The Modeling Framework

Preliminaries

I focus on two-agent interactions, purely for expositional simplicity. Let X be a set of *external states*. For each player $i = 1, 2$, let D_i be a set of pairs $d_i = (n_i, r_i)$ representing the player's *data access*, where n_i and r_i stand for his *news access* and *archival access* (to be defined explicitly below). Let A_i be the set of player i's feasible *actions*, and let Z be a set of *outcomes*. The

[4] In this example, θ has no direct payoff relevance, and it pins down players' archival access. These extreme assumptions serve the example's illustrative role. In a more realistic specification, θ would affect driving time, and the fraction of drivers with an app would vary more modestly with θ.

sets X and Z can be product sets, such that $x \in X$ or $z \in Z$ represent realizations of multiple variables. All sets are finite, for convenience.

I refer to $\Omega = X \times D_1 \times D_2 \times A_1 \times A_2 \times Z$ as the set of *states of the world*. Enumerate the variables that constitute states, such that $\omega \in \Omega$ can be written as $\omega = (\omega_1, ..., \omega_L)$. For every $B \subset \{1, ..., L\}$, denote $\omega_B = (\omega_l)_{l \in B}$ and $\Omega_B = \times_{l \in B} \Omega_l$. A state of the world resolves all uncertainty, covering exogenous variables (x, d_1, d_2) as well as endogenous ones (a_1, a_2, z). In applications, some state components may be degenerate (and therefore omittable). I refer to $\{1, ..., L\}$ as the set of variable *labels*, and often use the boldface notation $\mathbf{y}$ to indicate the label (or labels) of the variable (or collection of variables) y.

Impose the following structure on the notion of data access:

- $n_i \subseteq \{\mathbf{x}, \mathbf{d}_1, \mathbf{d}_2\}$. That is, player i's news access n_i is a subset of the set of labels of x, d_1, d_2. It defines the exogenous variables whose realization the player learns before taking an action. Assume further that $\mathbf{d}_i \subseteq n_i$ - i.e., the player is always informed of his own data access.[5] I refer to ω_{n_i} (the realization of the variables to which player i has news access) as player i's *news information*.

- $r_i \subseteq 2^{\{1,...,L\}}$. Player i's archival access r_i is defined by a collection of subsets of variables. As we will see below, the meaning of $B \in r_i$ is that player i learns the objective (steady-state) distribution over Ω_B. Denote $U(r_i) = \cup_{B \in r_i} B$.

Endow each player i with a von-Neumann-Morgenstern (vNM) utility function $u_i : \Omega \to \mathbb{R}$. Assume that u_i *is measurable with respect to* $(\omega_{U(r_i)}, a_i)$

[5] This assumption is behaviorally meaningful, but I do not regard it as a critical feature of the formalism. I introduce it mainly to simplify notation later on.

for any possible realization of r_i. That is, the player always has archival access to variables that are necessary for defining his payoffs.[6]

Let $\mu \in \Delta(X \times D_1 \times D_2)$ be an objective distribution over the exogenous variables. Player i's strategy given his news access n_i is $\sigma_{i,n_i} : \Omega_{n_i} \to \Delta(A_i)$. I refer to $\sigma_i = (\sigma_{i,n_i})_{n_i}$ as player i's *strategy*. As in standard simultaneous-move games, players' strategies are independent conditional on their news information. Let $h = (h(z \mid x, d_1, d_2, a_1, a_2))$ be the outcome distribution conditional on all other variables.

Define $p \in \Delta(\Omega)$ as the joint distribution over states of the world induced by these components - i.e., $p = \mu \circ \sigma_1 \circ \sigma_2 \circ h$. This distribution combines exogenous components (μ, h) and *endogenous* ones (σ_1, σ_2). I refer to p as a *steady-state distribution*. Unless indicated otherwise, p has full support. Let p^B denote the marginal of p on Ω_B. In a convenient abuse of notation, I will omit the superscript B when it is obvious from the context. For example, $p(x, a_1, a_2)$ is the marginal probability of (x, a_1, a_2) according to p.

The archival-access variable r_i represents player i's piecemeal knowledge of p. Each $B \in r_i$ means that player i learns p^B. I refer to the collection $(p^B)_{B \in r_i}$ as player i's *archival information*. When $\{1, ..., L\} \in r_i$, player i has complete archival information, as he simply learns p. In contrast, when r_i consists of strict subsets of $\{1, ..., L\}$, player i's archival information is incomplete.

Belief formation

To make a decision, player i forms a probabilistic belief based on his information. In the standard model of Bayesian games, information is synonymous with news information, and the player forms his belief in a single step: Bayesian updating of p conditional on his information. In the present model, the player has two kinds of information, and so he forms his conditional prob-

[6]It is possible for player i's archival access to take two possible values r_i, r'_i such that $U(r_i) \neq U(r'_i)$, and yet u_i will be measurable with respect to both $(\omega_{U(r_i)}, a_i)$ and $(\omega_{U(r'_i)}, a_i)$. The reason is that some variables may be redundant for pinning down the player's payoffs. Note also that $\omega_{U(r_i)}$ may already include a_i.

abilistic assessment in two stages, where each stage makes use of a different kind of information:

Stage one involves *maximum-entropy extrapolation from archival information*: The player forms the unconditional belief $p_{r_i} \in \Delta(\Omega_{U(r_i)})$ that solves

$$\max_{q \in \Delta(\Omega_{U(r_i)})} \left[- \sum_{y \in \Omega_{U(r_i)}} q(y) \ln(q(y)) \right] \tag{2}$$
$$\text{s.t. } q^B \equiv p^B \text{ for every } B \in r_i$$

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

Stage two involves *conditioning on the player's news information* and his action, according to conventional Bayesian updating. The player's conditional belief over $\Omega_{U(r_i)}$ is thus $(p_{R_i}(\omega_{U(r_i)} \mid \omega_{n_i}, a_i))$.[7]

Thus, each component of the player's information is associated with a particular *operation* that he performs on the objective distribution p. The first (second) stage involves *extrapolation* (*conditioning*); the player's archival (news) information determines what he extrapolates from (conditions on).

The procedure's second stage utilizes the canonical rule of Bayesian updating. By comparison, there is no "canonical" extrapolation rule for the first stage. Nevertheless, there is a common intuition that extrapolating a belief from partial data should follow some *parsimony* criterion. Maximum entropy (which originates from statistical physics and has a rich tradition in data analysis since Jaynes (1957)) offers one way to systematize this idea,

[7] When player i lacks archival access to some variable in (ω_{n_i}, a_i), we can safely remove it from the list of conditioned variables.

by formalizing the idea that the player forms a belief that has as little structure as possible subject to being consistent with observed correlations. We will see later that the criterion subsumes a number of precedents from the literature as special cases.

Equilibrium

Having defined players' beliefs as a function of their information, we are ready to introduce the notion of equilibrium, which is a standard trembling-hand perfection concept. The need for trembles arises because otherwise, players' second-stage, conditional belief may involve conditioning on null events. The trembles are conventionally interpreted as blind experimentation.

Definition 1 *Fix $\varepsilon > 0$, as well as μ and h (the exogenous components of p). A profile of full-support strategies (σ_1, σ_2) is an ε-equilibrium if for every $i = 1, 2$ and every n_i, ω_{n_i}, a_i for which $\sigma_{i,n_i}(a_i \mid \omega_{n_i}) > \varepsilon$,*

$$a_i \in \arg\max_{a_i' \in A_i} \sum_{y \in \Omega_{U(r_i)}} p_{r_i}(y \mid \omega_{n_i}, a_i') u_i(y, a_i')$$

A profile of strategies (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 $\varepsilon \to 0$.

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.

More on the interpretation of p and archival access

I interpret p as a representation of a long historical record of similar interactions by many agents who assumed the players' roles. Each agent moves *once* against this historical background.

I offer two concrete images for the notion of partial archival access. First, the player can receive statistical data about p from multiple sources, each recording a different collection of variables. Each of player i's data sources corresponds to a distinct subset $B \in r_i$. E.g., in a common-value auction, one data source may describe how bidders' signals vary with the object's true value, while another source describes how bidders' behavior varies with their signal.

Second, we can visualize p as an infinitely long *spreadsheet*, where $\{1, ..., L\}$ is the set of column titles (or "data fields") and rows correspond to independent draws from p. When player i obtains the spreadsheet, some of its values are missing according to some independent random process. In particular, for each spreadsheet row, the set of variables whose values are *not* missing is some $B \in r_i$.

In a number of our examples, the label $\mathbf{r}_j$ will belong to some $B \in r_i$. This means that player i has archival access to player j's archival access. For such a configuration to be sensible, the variable r_j should be observable in principle. Indeed, in most of the examples in this paper, I give r_i a *physical* interpretation that makes it possible for us to imagine it being an observable variable. For instance, in Section 2, r_i indicated whether player i has a navigation app.

Closed Forms of p_r

In many cases of interest, the maximum-entropy extension of a given collection of marginal distributions is tractable. The simplest case is where r_i consists of a single subset $\{B\}$, $B \subset \{1, ..., L\}$. In this case, $p_{r_i} = p^B$. Though apparently trivial, this case can have significant behavioral implications, as we saw in Section 2.

For a slightly more complicated example, suppose that r_i consists of *two* subsets, $r_i = \{B, C\}$. This means that player i learns p^B and p^C. The maximum-entropy extension of these marginals is

$$p_{r_i}(\omega_B, \omega_C) = p(\omega_B)p(\omega_{C-B} \mid \omega_{C\cap B})$$

This formula transparently maximizes statistical independence subject to the known correlations. Most of the examples in this paper make use of this special case. In Section 4.3, I present a generalization for cases in which r_i consists of more than two subsets. Precedents in the literature for some of these formulas will be discussed in Section 6.

Do players understand they are playing a game?

I presented the objective environment in approximately game-theoretic terms. The description is unconventional because of the definition of the state of the world, which served my definition of archival access. But while the game-theoretic description may have been non-standard, it followed the familiar causal structure: exogenous variables are realized first; then players act independently, conditioning their behavior on their information; and finally, the outcome is realized. However, this objective causal structure is *not* essential for our description of players' belief formation and behavior. As the "spreadsheet" metaphor suggests, the underlying assumption is that players lack an understanding of the rules of the game and its causal structure. They need not even understand they are playing a game. Their perception of the situation is derived entirely from the partial statistical datasets at their disposal. From this point of view, the objective causal structure of p is superfluous, and presented here only for convenience.

Relation to the standard Harsanyi type space

My occasional reference to d_i as player i's "type" naturally brings to mind the standard Harsanyi type space, and raises the question of whether it can embed the "type space" of this paper. The answer is generally negative, because the maximum-entropy extension p_{r_i} need not satisfy Bayes plausibility - i.e., it is possible that

$$\sum_{\omega_B} p(\omega_B) p_{r_i}(\omega_C \mid \omega_B) \neq p(\omega_C)$$

for some p and ω and some subsets $B, C \subset U(r)$. That is, the expected subjective posterior probability of ω_C (where the expectation is taken with respect to the objective distribution) does not always coincide with the objective prior probability of ω_C. It can be easily verified that some of the examples in this paper have this feature. For this reason, players' beliefs in this formalism cannot in general be reproduced by a conventional Harsanyi type space, where beliefs *do* obey Bayes plausibility.

At a higher level, the two typologies are quite different in spirit. In the standard Harsanyi framework, players have well-defined (common or subjective) prior beliefs over the state space (whether or not the definition of a state includes endogenous variables). In contrast, the starting point of the present formalism is that players sometimes have a *piecemeal* perception of the (objective) distribution. If we try to represent this partial perception by a Harsanyi type in some extended state space, then we simply shift the problem to another level because we need to define players' prior beliefs over that extended state space.

4 Examples

This section presents examples that demonstrate the formalism's expressive scope. In all examples, I hold players' news access constant - i.e., the set of variables whose realization a given player learns before taking an action is fixed. This enables me to omit n_1 and n_2 from the definition of a state of the world, and focus on the formalism's novel aspect: random archival access.

4.1 "Market Savvy": Correlation between Archival Access and other Player Characteristics

What determines market agents' relative performance? Do "savvier" agents earn higher profits?[8] The following example highlights the non-trivial role

[8] Piccione and Rubinstein (2003) and Eyster and Piccione (2013) explored this question in the context of competitive asset market models. Applying this paper's terminology retroactively, traders in the Eyster-Piccione model are characterized by incomplete archive-

of the correlation between two aspects of "market savvy": market players' ability to exploit market opportunities and their ability to assess the extent to which their opponent chases (and therefore dissipates) the same opportunities.

In the example, the two players are entrepreneurs who consider entering a certain field of activity. The variable $\theta \in \{0,1\}$ indicates whether there is a viable market in this field. The variable $s_i \in \{0,1\}$ indicates whether player i has the relevant technology. The variable $a_i \in \{0,1\}$ indicates whether the player enters the field. When the trader lacks the technology ($s_i = 0$), he is forced to remain inactive ($a_i = 0$); he only has a choice when $s_i = 1$.

Player i's payoff is $u_i(\theta, a_1, a_2) = \theta a_i(3 - 2a_i - 2a_j)$. The story behind this specification is as follows. When $\theta = 0$, entrepreneurs are unable to turn a profit in this field. When $\theta = 1$, the market is viable, but has no room for two active players. If only one player enters, he earns a profit of 1. If both players enter, each loses 1. The structure of u_i implies that players' beliefs regarding θ are irrelevant; what matters for player i's decision is his prediction of a_j *conditional on* $\theta = 1$. In particular, he strictly prefers to play $a_i = 1$ (given $s_i = 1$) if and only if $p_{r_i}(a_j = 1 \mid \theta = 1) < \frac{1}{2}$.

Players' news access is always $n_i = \{\mathbf{s}_i, \mathbf{d}_i\}$ - i.e., each player always knows his technology and data access. Player i's archival access can take two values, referred to by the shorthand notation 0 and 1 and given explicitly as follows: $r_i = 0$ signifies $\{\{\boldsymbol{\theta}, \mathbf{s}_j\}, \{\mathbf{s}_j, \mathbf{a}_j\}\}$, and $r_i = 1$ signifies $\{\{\boldsymbol{\theta}, \mathbf{a}_j\}\}$. That is, $r_i = 1$ means that player i fully grasps how player j's actions correlate with θ. In contrast, $r_i = 0$ means that the player lacks direct evidence regarding this correlation: he only learns how player j's technology varies with market conditions, and how his actions vary with his technology.

The prior distribution over $\theta, r_1, s_1, r_2, s_2$ is:

(*i*) $p(\theta = 1) = \frac{1}{2}$.

(*ii*) $p(r_i = 1) = \frac{1}{2}$, independently of θ and of r_j.

information, given by a single set of variable labels, which differs across trader types.

(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 $q_1 \neq q_0$ and $q_1, q_0 \geq \frac{1}{2}$.

The q parameters capture the correlation between market viability and players' technology. When s_i is highly correlated with θ, the player is a "*savvy*" entrepreneur in the sense that he has good "market timing". If the player faced no competition, higher correlation between s_i and θ would lead to larger expected profits. However, "savviness" has another dimension, given by the players' archival access: A player with $r = 1$ has better archival knowledge of his opponent's behavior. The relation between q_1 and q_0 determines the correlation between the two aspects of market savvy. When $q_1 > q_0$ ($q_1 < q_0$), the correlation is positive (negative): a player with better archival access has better (worse) market timing.

I analyze symmetric equilibria. Let $\alpha_r = p(a = 1 \mid s = 1, r)$ represent the equilibrium strategy of a player with archival access $r \in \{0, 1\}$.

Rational-expectations benchmark

Consider a benchmark model in which every realization of r_i induces rational expectations - i.e., $p_{r_i}(a_j \mid \theta) \equiv p(a_j \mid \theta)$. The only possible difference between realizations of r is thus the value of q_r. This reduces the model to a conventional game with incomplete information. Any symmetric Nash equilibrium generates *zero profits* for both players, independently of their value of r, such that $p(a_i = 1 \mid \theta = 1) = \frac{1}{2}$ for every $i = 1, 2$. In particular, there is a symmetric Nash equilibrium in which players' behavior is

$$\alpha_r = \frac{1}{q_0 + q_1} \tag{3}$$

for all r. I refer to it as the *type-independent equilibrium*. There are also type-dependent equilibria; both $\alpha_1 > \alpha_0$ and $\alpha_1 < \alpha_0$ are consistent with equilibrium.

Let us now switch back from this benchmark to our original specification. The following result characterizes the set of symmetric equilibria.

Proposition 1 *(i) When* $q_1 > q_0$*, the unique symmetric equilibrium is the type-independent equilibrium.*

(ii) When $q_1 < q_0$*, there is exactly one symmetric equilibrium in addition to the type-independent equilibrium: players with* $r = 1$ *play* $\alpha_1 = 1$ *and earn a positive profit; players with* $r = 0$ *play* $\alpha_0 = (2 - q_0 - q_1)/(q_0 + q_1)$ *and earn zero profits.*

Proof. Let us first construct players' beliefs as a function of their archival access. When $r_i = 1$, player i's archival information fully documents the joint distribution of θ and a_j. Therefore, his conditional prediction is consistent with rational expectations:

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

When $r_i = 0$, player i's first-stage belief over θ, s_j, a_j can be written as

$$p_{r_i=0}(\theta, s_j, a_j) = p(\theta)p(s_j \mid \theta)p(a_j \mid s_j)$$

This induces the following perceived conditional distribution:

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

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

$$p_{r_i=0}(a_j = 1 \mid \theta = 1) = p(s_j = 1 \mid \theta = 1)p(a_j = 1 \mid s_j = 1) \tag{4}$$

The following elaboration of these formulas 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, such that $p_{r_i=0}(a_j = 1 \mid \theta = 1)$ can be written as

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

Consider a symmetric equilibrium. Assume $\alpha_0 = \alpha_1 = 0$ - i.e., $p(a_j = 1 \mid s_j = 1, r_j) = 0$ for every r_j. By assumption, $p(a_j = 1 \mid s_j = 0, r_j) = 0$ for every r_j. It follows that $p_{r_i}(a_j = 1 \mid \theta = 1) = 0$ for any realization of r_i. Therefore, it would be profitable for any player i to deviate to $a_i = 1$ when $s_i = 1$, a contradiction. Therefore, $\alpha_0 > 0$ or $\alpha_1 > 0$. Using similar reasoning, we can show that $\alpha_0 < 1$ or $\alpha_1 < 1$. As observed earlier, player i chooses $a_i = 1$ with positive probability only if $p_{r_i}(a_j = 1 \mid \theta = 1) \leq \frac{1}{2}$. Therefore, $\alpha_1 > \alpha_0$ only if

$$p_{r_i=1}(a_j = 1 \mid \theta = 1) \leq p_{r_i=0}(a_j = 1 \mid \theta = 1) \tag{5}$$

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

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

whereas

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

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, (5) becomes

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

which is equivalent to

$$\alpha_1(q_1 - q_0) \leq \alpha_0(q_1 - q_0) \tag{6}$$

Suppose $q_1 > q_0$. Then, this inequality contradicts the inequality $\alpha_1 > \alpha_0$. A similar contradiction is obtained for $\alpha_1 < \alpha_0$. It follows that when $q_1 > q_0$, we must have $\alpha_1 = \alpha_0 \in (0,1)$ in equilibrium. That is, the only possible symmetric equilibrium is the one given by (3), such that players earn zero profits, regardless of their archival access.

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 $\alpha_1 \in (0,1)$, players with $r = 1$ are indifferent between the actions. This indifference means $p(a_j = 1 \mid \theta = 1) = \frac{1}{2}$, such that all players earn zero objective profits, independently of their archival access. If $\alpha_1 > \alpha_0$, (5) must hold with strict inequality, such that players with $r = 0$ strictly prefer $a = 0$, hence $\alpha_0 = 0$. But this means that $p(a_j = 1 \mid \theta = 1) < \frac{1}{2}$, contradicting the indifference of players with $r = 1$. Likewise, if $\alpha_1 < \alpha_0$, (5) must hold with strict inequality, such that players with $r = 0$ strictly prefer $a = 1$, hence $\alpha_0 = 1$. But this means that $p(a_j = 1 \mid \theta = 1) > \frac{1}{2}$, contradicting the indifference of traders with $r = 1$. It follows that when $q_1 < q_0$, we must have $\alpha_1 \in \{0, 1\}$. If $\alpha_1 = 0$, we obtain a contradiction with (6). Therefore, the only remaining possibility is $\alpha_1 = 1$ and $\alpha_0 \in (0, 1)$, such that players with $r = 0$ are indifferent between $a = 0$ and $a = 1$ - i.e., $p_{r_i=0}(a_j = 1 \mid \theta = 1) = \frac{1}{2}$. This equation yields the solution for α_0. Plugging this value in $p_{r=1}(a_j = 1 \mid \theta = 1)$, we can verify that players with $r = 1$ earn positive profits, consistent with the assumption that $\alpha_1 = 1$. ■

Thus, if $q_1 > q_0$ - i.e., a player with superior archival access also has better market timing - players' equilibrium behavior must be independent of their archival access and they all must earn zero profits. The savvy player's "double advantage" ends up having no effect on his equilibrium market performance.

The reasoning behind this result is roughly as follows. Suppose that in equilibrium, players with richer archival access have a higher propensity to enter conditional on $s = 1$. Thus, r_j is positively correlated with both s_j and a_j. A player i with $r_i = 0$ perceives the positive correlation between θ and a_j only through these variables' correlation with s_j. This leads him to underestimate the correlation between θ and a_j, and by implication player j's propensity to enter conditional on $\theta = 1$. But this should make him *more* eager to enter than a player i with $r_i = 1$, a contradiction. A similar contradiction is reached if we assume that players' eagerness to enter decreases with r.

When $q_1 < q_0$, we can sustain exactly one symmetric equilibrium with non-zero profits: players with $r = 1$ always choose to be active when possible, whereas players with $r = 0$ only do so with some probability. Both player types earn positive objective profits, yet players with $r = 0$ overestimate the competition and therefore wrongly predict zero profits.

When $q_1 = q_0$, s_j is independent of r_j, hence r_j does not confound the correlation between s_j and a_j. Players with $r = 0$ and $r = 1$ must form the same beliefs. As a result, the model is reduced to the rational-expectations benchmark and players earn zero profits in equilibrium.

4.2 "Market Intelligence" and Collusive Outcomes

This sub-section illustrates the formalism's ability to capture situations in which one player has partial archival information regarding his opponent's archival access. The setting is a stylized model of market competition. Players are firms that adapt the quality of their products to the size of consumer demand. Each firm may have some "market intelligence" regarding the competitive strategy of its rival - namely, how the rival firm's behavior varies with underlying demand or with its level of market intelligence. This kind of market intelligence can be described in terms of archival access. The question is how the firms' random archival access affects their ability to sustain non-competitive behavior.

Specifically, each firm's action a_i can take two values, L and H, representing basic and premium product quality. The cost of serving basic and premium products is 0 and 1, respectively. When both firms choose the same product quality, each gets a market share of 50%; if exactly one firm offers a premium product, it gets 100% of the market. The size of consumer demand is $x \in \{b, g\}$; the two values of x are equally likely. This description induces the following payoff matrix:

$a_1 \backslash a_2$	L	H
L	$\frac{x}{2}, \frac{x}{2}$	$0, x-1$
H	$x-1, 0$	$\frac{x-1}{2}, \frac{x-1}{2}$

I impose the following restrictions on the parameters b and g:

$$g > 2 \geq b > \frac{3}{2}$$

These restrictions imply that H is a strictly dominant action when $x = g$, whereas both (L, L) and (H, H) are Nash equilibria when $x = b$. Furthermore, H is the risk-dominant action when $x = b$, because the following inequality is satisfied:

$$\frac{1}{2} \cdot (b-1) + \frac{1}{2} \cdot \frac{b-1}{2} > \frac{1}{2} \cdot \frac{b}{2} + \frac{1}{2} \cdot 0$$

This type of payoff matrix is familiar from the "global games" literature (Rubinstein (1989), Carlsson and van Demme (1993), Morris and Shin (1998)). I refer to (H, H) and (L, L) as the "competitive" and "collusive" outcomes, respectively. This payoff matrix will also serve us in Sections 4.2.1 and 4.2.2.

Each firm's archival access r_i takes two possible values, given the shorthand notation 0 and 1, where $r_i = 0$ represents $\{\{\mathbf{x}\}, \{\mathbf{a}_j\}\}$ (poor archival access) and $r_i = 1$ represents $\{\{\mathbf{x}, \mathbf{r}_j, \mathbf{a}_j\}\}$ (rich archival access). To see why r_j can be viewed as an observable variable (such that statistical data about its realizations can be available), note that players' archival access can be determined by their experience or whether consult a market specialist.

When $r_i = 0$, firm i learns the marginal distributions of consumer demand and the opponent's action, but has no data about their correlation. To use the terms of Eyster and Rabin (2005), $r_i = 0$ represents a "fully cursed" player.[9] When $r_i = 1$, firm i learns how firm j's product quality varies with consumer demand as well as with j's archival access. However, firm i never learns how firm j conditions its behavior on *firm i's own* archival access. In other words, the firm learns behavioral patterns by different types of its opponent, yet it lacks data about how its opponent behaves against different opponent types.

Assume $p(r_i = 1) = \beta$ for each player i. I allow r_1 and r_2 to be correlated. In particular, $p(r_i = 1 \mid r_j = 1) = \gamma$ for both $i = 1, 2$, $j \neq i$. A low value of γ captures situations in which high-quality market intelligence is like an exclusive product: when one firm enjoys it, the other firm is likely to lack it.

[9] Ettinger and Jehiel (2010) interpreted the failure to perceive any correlation between a player's behavior and the underlying state of Nature in terms of the so-called *Fundamental Attribution Error.*

Firms' news access is fixed: $n_1 = n_2 = \{\mathbf{x}, \mathbf{d}_1, \mathbf{d}_2\}$.[10]

Proposition 2 (i) *In any equilibrium, each player i plays $a_i = H$ whenever $x = g$ or $r_i = 0$. (ii) An equilibrium in which each player i chooses $a_i = L$ when $x = b$ and $r_i = 1$ exists, if and only if $\gamma \geq b - 1$.*

Proof. (i) Suppose $x = g$. Then, $a_i = H$ is a strictly dominant action for player i. Since p_{r_i} is not defined over a_i (unlike the example in Section 2), it is impossible for p_{r_i} to rationalize an objectively dominated action. Therefore, for any realization of r_i, each player will choose H when $x = g$ in equilibrium.

Now suppose $x = b$ and $r_i = 0$. By definition,

$$p_{r_i=0}(x, a_j) = p(x)p(a_j)$$

Then, player i will choose L over H only if

$$p(a_j = L) \cdot (b-1) + p(a_j = H) \cdot \frac{b-1}{2} \leq p(a_j = L) \cdot \frac{b}{2} + p(a_j = H) \cdot 0 \quad (7)$$

which can be simplified into $b \leq 1 + p(a_j = L)$. Since we saw that $a_j = H$ whenever $x = g$, it follows that $p(a_j = L) \leq \frac{1}{2}$. Then, (7) implies $b \leq \frac{3}{2}$, a contradiction.

(ii) Suppose $x = b$ and $r_i = 1$. Then,

$$\begin{aligned} p_{r_i=1}(a_j &= L \mid x = b, r_i = 1, r_j) = p(a_j = L \mid x = b, r_j) \\ &= \sum_{r_i'} p(r_i' \mid r_j) p(a_j = L \mid x = b, r_i', r_j) \end{aligned}$$

Let us calculate this expression. First, we have established that $a_j = H$ whenever $r_j = 0$. That is, $p_{r_i=1}(a_j = L \mid x, r_i = 1, r_j = 0) = 0$ for every x.

[10] It would be more natural to assume that high-quality market intelligence manifests itself in rich archival access *as well as* rich news access - e.g., when $r_i = 0$, firm i is only informed of x and r_i, whereas when $r_i = 1$ it is also informed of r_j. Since this would lead to the same results, I chose to make the simpler assumption that n_i is fixed.

Therefore, when $r_i = 1$ and $r_j = 0$, player i's best-reply is $a_i = H$ (for any value of x). Since the same argument applies to both players, it follows that $a_j = H$ when $r_i = 0$ and $r_j = 1$. Hence, $p_{r_i=1}(a_j = L \mid x = b, r_i = 1, r_j = 0) = 0$.

It follows that

$$\begin{aligned} p_{r_i=1}(a_j &= L \mid x = b, r_i = 1, r_j = 1) \\ &= p(r_i = 1 \mid r_j = 1)p(a_j = L \mid x = b, r_i = 1, r_j = 1) \end{aligned}$$

Therefore, we can sustain an equilibrium in which players choose L when $x = b$ and $r_1 = r_2 = 1$ only if

$$\gamma \cdot (b-1) + (1-\gamma) \cdot \frac{b-1}{2} \leq \gamma \cdot \frac{b}{2} + (1-\gamma) \cdot 0$$

which can be simplified into $\gamma \geq b - 1$. ■

Thus, the collusive outcome can be sustained in equilibrium when $x = b$ and $r_1 = r_2 = 1$, but only if $\gamma \geq b-1$. In particular, the condition fails when $\gamma \leq \frac{1}{2}$, given our restriction that $b > \frac{3}{2}$. Put differently, when high-quality market intelligence tends to be an exclusive product, firms cannot utilize it to sustain the collusive outcome - even when both firms happen to have it.

The logic is as follows. When $r_i = 1$, firm i understands how firm j conditions its behavior on r_j. However, since firm i lacks data about how firm j conditions its behavior on r_i, it implicitly averages over all values of r_i. When γ is low, $r_i = 0$ is the probable realization given $r_j = 1$. When $r_j = 1$, firm j understands that firm i chooses H when $r_i = 0$, and therefore firm j responds to $r_i = 0$ by playing H. It follows that when γ is low, firm i predicts that firm j is likely to play H even when $r_j = 1$. This in turn implies that firm i will prefer to play H as well.

The notion of imperfect archival access plays a key role in this result. Consider a benchmark model in which firms always have rational expectations, such that only their news access is random; specifically, high-quality market intelligence by firm i consists of learning the realization of r_j. In this benchmark model, the collusive outcome can be sustained in equilibrium when $x = b$ and $r_1 = r_2 = 1$. The reason is that firm i is always informed of its own type; under rational expectations, learning firm j's type enables firm i to make a precise prediction of j's behavior.

Does richer archival access necessarily facilitate collusion?
Consider the following variant on our example, in which $r_i = 1$ represents $\{\{\mathbf{x}, \mathbf{a}_j\}\}$. This is thinner than the archival access represented by $r_i = 1$ in Proposition 2; it means that firm i learns how j's behavior varies with consumer demand, but without any understanding of how it varies with players' archival access. Then,

$$\begin{aligned} p_{r_i=1}(a_j &= L \mid x = b, r_i = 1, r_j) = p(a_j = L \mid x = b) \\ &= \sum_{r_i} \sum_{r_j} p(r_i, r_j) p(a_j = L \mid x = b, r_i, r_j) \end{aligned}$$

Part (i) of Proposition 2 continues to hold - i.e., player j chooses H when $r_j = 0$. Also, since $p_{r_j=1}$ is not defined over r_i, player j's behavior is independent of r_i even when $r_j = 1$. Therefore,

$$p_{r_i=1}(a_j = L \mid x = b, r_i = 1, r_j) \leq p(r_j = 1) = \beta$$

It follows that player i will choose L when $x = b$ and $r_i = 1$ only if

$$\beta \cdot (b - 1) + (1 - \beta) \cdot \frac{b-1}{2} \leq \beta \cdot \frac{b}{2} + (1 - \beta) \cdot 0$$

which can be simplified into $\beta \geq b - 1$. This is the condition for sustaining the collusive outcome in equilibrium when $x = b$ and $r_1 = r_2 = 1$. When $\beta > \gamma$, this condition is *less* stringent than in Proposition 2. Therefore, richer archival access need not facilitate collusion.

4.3 Larger Archival-Access Sets

In Section 3, I observed that in some cases, maximum-entropy extension has simple closed forms. I now generalize these cases into a single principle.

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

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

Proposition 3 (Hajek et al. (1992)) *When r satisfies RIP, the maximum-entropy extension of $(p^B)_{B \in R}$ is given by*

$$p_r(\omega_{U(r)}) = \prod_{B^1,...,B^m} p(\omega_{B^k - (\cup_{j<k} B^j)} \mid \omega_{B^k \cap (\cup_{j<k} B^j)}) \tag{8}$$

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

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

$$\begin{aligned} p_r(\omega_1, \omega_2, \omega_3, \omega_4) &= p(\omega_1, \omega_2) p(\omega_3 \mid \omega_2) p(\omega_4 \mid \omega_3) \\ &= p(\omega_1) p(\omega_2 \mid \omega_1) p(\omega_3 \mid \omega_2) p(\omega_4 \mid \omega_3) \end{aligned} \tag{9}$$

Thus, RIP allows us to write p_r as a factorization of $p_r(\omega_{U(r)})$ into marginal and conditional distributions. The factorization makes it manifest that the player's belief maximizes statistical independence subject to the known correlations. Moreover, the factorization has a *causal* interpretation. For instance, (9) looks as if p_r is consistent with the causal chain $1 \to 2 \to 3 \to 4$.

This is a general property of (8). This formula can be rewritten as a factorization of p according to a *directed acyclic graph* (DAG) whose set of nodes is $U(r)$, such that r is the set of maximal cliques in the graph's non-directed version. Indeed, RIP is a familiar concept in the literature on graphical probabilistic models (see Cowell et al. (1999)). I refer the reader to Spiegler (2017) for more details.

Armed with the more general formula (8), let us now revisit the payoff matrix presented at the beginning of Section 4, and examine a different type space. Each firm i's news access is $n_i = (\mathbf{x}, \mathbf{d}_i)$. Thus, unlike the previous specification in this section, here one firm is not informed of its opponent's data access. Since n_i is fixed, we can omit it from the description of the state of the world, which can be written as $\omega = (x, r_1, r_2, a_1, a_2)$, The distribution over x is the same as before.

As before, r_i gets two values, given the shorthand notation 0 and 1. Assume that r_1 and r_2 are distributed *uniformly* and independently of x. Moreover, they are *positively* correlated: $p(r_i = r_j \mid r_j) = \delta \in (\frac{1}{2}, 1)$ for every $r_j = 0, 1$. Firms' archival access is given explicitly as follows:

$$\begin{aligned} r_i &= 0 : \{\{\mathbf{x}, \mathbf{a}_j\}\} \\ r_i &= 1 : \{\{\mathbf{x}, \mathbf{a}_j\}, \{\mathbf{r}_i, \mathbf{r}_j\}, \{\mathbf{r}_j, \mathbf{a}_j\}\} \end{aligned}$$

Thus, firm i has random archival access to firm j's archival access. Specifically, $r_i = 0$ means that firm i only learns how its opponent's action correlates with x, whereas $r_i = 1$ means that firm i also learns the pairwise correlations of firm j's archival access with its action and with firm i's own archival access. Thus, $r_i = 1$ represents richer archival access and therefore indicates higher-quality "market intelligence". Both realizations of r_i satisfy RIP, such that

$$
\begin{aligned}
p_{r_i=0}(x, a_j) &= p(x)p(a_j \mid x) \\
p_{r_i=1}(x, a_j, r_i, r_j) &= p(x)p(a_j \mid x)p(r_j \mid a_j)p(r_i \mid r_j)
\end{aligned}
$$

In this example, there are equilibria in which firms' behavior is independent of their archival access - and therefore coincides with the rational-expectations benchmark. The reason is that if a_j is independent of r_j, the realizations $r_i = 1$ and $r_i = 0$ both induce $p_{r_i}(a_j \mid x, r_i, a_i) \equiv p(a_j \mid x)$, and therefore there is no reason for firm i to vary its action with r_i.

Let us now examine whether there are equilibria that allow firms' actions to vary with their archival access.

Proposition 4 *There is a symmetric pure-strategy equilibrium in which firms' actions vary with their archival access, if and only if*

$$
\delta \geq \frac{2b - 2}{5 - 2b} \tag{10}
$$

In this equilibrium, firms choose H whenever $x = g$; and when $x = b$, each firm i chooses L if and only if $r_i = 1$.

Proof. In any equilibrium, any firm i chooses $a_i = H$ whenever $x = g$, independently of R_i. The reasoning is the same as in Proposition 2 and therefore omitted here. Let us now derive firm i's conditional belief over a_j conditional on $x = b$ and each of the two realizations of r_i:

$$
p_{r_i=0}(a_j = L \mid x = b, r_i = 0) = p(a_j = L \mid x = b)
$$

and

$$\begin{aligned} p_{r_i=1}(a_j &= L \mid x=b, r_i=1) = \frac{p_{r_i=1}(x=b, a_j=L, r_i=1)}{\sum_{a_j'} p_{r_i=1}(x=b, a_j', r_i=1)} \\ &= \frac{p(x=b)p(a_j=L \mid x=b)\sum_{r_j} p(r_j \mid a_j=L)p(r_i=1 \mid r_j)}{p(x=b)\sum_{a_j'} p(a_j' \mid x=b)\sum_{r_j} p(r_j \mid a_j')p(r_i=1 \mid r_j)} \end{aligned}$$

which can be rewritten as

$$\left(1 + \frac{p(a_j=H \mid x=b)[p(r_j=1 \mid a_j=H)\cdot\delta + p(r_j=0 \mid a_j=H)\cdot(1-\delta)]}{p(a_j=L \mid x=b)[p(r_j=1 \mid a_j=L)\cdot\delta + p(r_j=0 \mid a_j=L)\cdot(1-\delta)]}\right)^{-1}$$

Suppose that when $x=b$, firms vary their (symmetric, pure-strategy) equilibrium action with their archival access. Then, $p(a_j=L \mid x=b)=\frac{1}{2}$. Because H is the risk-dominant action when $x=b$, it follows that when $x=b$ and $r_i=0$, firm i's best-reply is $a_i=H$. Thus, if we wish to construct a symmetric pure-strategy equilibrium in which a_j varies with r_j, it must be the case that $p(a_j=L \mid x=b, r_j)=r_j$. We can now calculate $p_{r_i=1}(a_j=L \mid x=b, r_i=1)$, by plugging the terms

$$\begin{aligned} p(a_j &= L \mid x=b) = \frac{1}{2} \\ p(r_j &= 1 \mid a_j=L) = 1 \end{aligned}$$

and

$$\begin{aligned} p(r_j &= 1 \mid a_j=H) = \frac{p(r_j=1, a_j=H)}{p(a_j=H)} \\ &= \frac{p(r_j=1)p(x=g)}{p(x=g)+p(x=b)p(r_j=0)} = \frac{\frac{1}{2}\cdot\frac{1}{2}}{\frac{1}{2}+\frac{1}{2}\cdot\frac{1}{2}} = \frac{1}{3} \end{aligned}$$

Thus,

$$p_{r_i=1}(a_j=L \mid x=b, r_i=1) = \left(1 + \frac{\frac{1}{2}\cdot[\frac{1}{3}\cdot\delta+\frac{2}{3}\cdot(1-\delta)]}{\frac{1}{2}\cdot[1\cdot\delta+0\cdot(1-\delta)]}\right)^{-1} = \frac{3\delta}{2+2\delta}$$

Therefore, the condition that $a_i = L$ is a best-reply when $x = b$ and $r_i = 1$ can be written as follows:

$$\frac{3\delta}{2+2\delta} \cdot \frac{b}{2} \geq \frac{3\delta}{2+2\delta} \cdot (b-1) + \frac{2-\delta}{2+2\delta} \cdot \frac{b-1}{2}$$

which can be simplified into (10). It follows that this inequality is necessary and sufficient for the existence of a symmetric pure-strategy in which firms sometimes vary their action with their archival access. ■

Once again, we see how random archival access can be an impediment to a collusive outcome. Even when the richness of firms' archival access is positively correlated, successful coordination depends on this correlation being strong enough.

A new effect arises here because $r_i = 1$ consists of *three* interlocking subsets, which enable firm i to learn the pairwise correlations of a_j with x and r_j, without enabling it to learn the *joint* correlation of a_j with x, r_j. This limitation distorts the way the firm updates its belief over a_j. Specifically, when firm i learns that $x = b$ and $r_i = 1$, it effectively regards these realizations as *conditionally independent signals* of a_j. This is an erroneous belief, because a_j as determined by the joint realization of x and r_j (the latter being positively correlated with r_i). This error attenuates firm i's confidence that player j will play L when $x = b$ and $r_i = 1$:

$$p_{r_i=1}(a_j = L \mid x = b, r_i = 1) = \frac{3\delta}{2+2\delta} < \delta = p(a_j = L \mid x = b, r_i = 1)$$

The belief distortion thus makes it harder for firms to sustain a collusive outcome in equilibrium. By comparison, consider the alternative specification in which $r_i = 1$ represents the collection $\{\{\mathbf{x}, \mathbf{r}_j, \mathbf{a}_j\}, \{\mathbf{r}_i, \mathbf{r}_j\}\}$, such that $p_{r_i=1}(a_j \mid x, r_i = 1)$ coincides with rational expectations. Then, the equilibrium I constructed in Proposition 4 is sustainable whenever $\delta \geq b - 1$, which is a less stringent condition than (10).

4.4 Hierarchical Archival Access

The representation of a state of the world in terms of a collection of variables is fundamental to the modeling framework. Furthermore, unlike the standard Harsanyi model, replacing a collection of variables with a single variable is not innocuous. For instance, let $\theta = (\theta^1, ..., \theta^K)$ be a state of Nature. If we collapse θ into a single variable, we cannot capture situations in which one player lacks archival access to one of the components of θ. Perhaps the most interesting case of this effect is where the archival-access variable r_i itself corresponds to a collection of variables. This sub-section develops this idea.

In some of the examples in this paper, players had archival access to archival-access variables. The formalism's capacity for such cross-references is one of its prime virtues - in rough analogy to the Harsanyi formalism's ability to describe one player's news information regarding another player's news information. Inspired by this analogy (despite the very different interpretation), I am led to think of *hierarchical constructions* of such inter-dependence.

The starting point of a hierarchical definition of players' archival access is a collection of *basic* variables. Let $B \subset \{1, ..., L\}$ be the set of labels of the basic variables. These would include variables that define the external state, players' data access, players' actions and the outcome. For each player i, there is a collection of variables $r_i^1, ..., r_i^m$, $m \geq 2$, where $r_i^1 \subset 2^B$, and for every $k = 2, ..., m$,

$$r_i^k \subset 2^{B \cup \{\mathbf{r}_i^h, \mathbf{r}_j^h\}_{h=1,...,k-1}}$$

In addition, every element in r_i^k, $k \geq 2$, includes $\mathbf{r}_i^{k-1}$ or $\mathbf{r}_j^{k-1}$. Define $r_i = \cup_{k=1,...,m} r_i^k$ and regard it as a distinct variable.

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

The following is a simple example of hierarchically defined archival access in the context of the "global game" that serves as a running example in this section. The distribution over x and the payoff matrix are exactly the same, and only firms' data access is modified. Firms have common archival access, given by the list $(r^1, r^2, ...)$ and distributed independently of x. Players' news access is fixed: $n_1 = n_2 = \{\mathbf{x}, \mathbf{n}_1, \mathbf{n}_2, \mathbf{r}^1, \mathbf{r}^2, ...\}$. The basic variables are x, a_1, a_2. For each k, r^k takes two values, given the shorthand notation 0 and 1 and defined explicitly as follows:

k	$r^k = 0$	$r^k = 1$
1	$\{\{\mathbf{x}\}, \{\mathbf{a}_1\}, \{\mathbf{a}_2\}\}$	$\{\{\mathbf{x}, \mathbf{a}_1, \mathbf{a}_2\}\}$
2	$\emptyset$	$\{\{\mathbf{x}, \mathbf{a}_1, \mathbf{a}_2, \mathbf{r}^1\}\}$
3	$\emptyset$	$\{\{\mathbf{x}, \mathbf{a}_1, \mathbf{a}_2, \mathbf{r}^1, \mathbf{r}^2\}\}$
$\vdots$	$\vdots$	$\vdots$
m	$\emptyset$	$\{\{\mathbf{x}, \mathbf{a}_1, \mathbf{a}_2, \mathbf{r}^1, ..., \mathbf{r}^{m-1}\}\}$
$\vdots$	$\vdots$	$\vdots$

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 give r a shorthand notation, namely the largest value of k for which $r^k = 1$. Specifically, let $p(r = k) = \gamma(1-\gamma)^k$ for every $k = 0, 1, ...$ Note that $r = k$ means that firms perceive actions as a function of $x, r^1, ..., r^{k-1}$.

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

Proof. The proof is by induction on k. As a first step, observe that by the same argument as in Proposition 2, $a_i = H$ whenever $x = g$ or $r = 0$. Suppose that we have shown that $a_i = H$ when $x = b$ and $r < k$, and consider the case of firm 1, say, when $x = b$ and $r = k$. The firm will find it optimal to play $a_1 = L$ only if

$$p_{r=k}(a_2 = L \mid x = b, r^1 = \cdots = r^k = 1, r^{k+1} = r^{k+2} = \cdots = 0, a_1 = 1) > \frac{1}{2}$$

The L.H.S of this inequality can be written as

$$\begin{aligned} p(a_2 &= L \mid x = b, r^1 = \cdots = r^{k-1} = 1) \\ &\leq \frac{p(r \geq k)}{p(r \geq k-1)} = 1 - \gamma < \frac{1}{2} \end{aligned}$$

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

The intuition for this result is as follows. When firms have $r = k$, they only perceive correlations between actions and archival access of order $k - 1$ and below. By the assumption that $\gamma > \frac{1}{2}$, firms are more likely to lack k^{th}-order archival access conditional on having $(k-1)^{th}$-order archival access. By the inductive step, they play $a = H$ in that case. It follows that when firms have $r = k$, they believe that the opponent is more likely to play $a = H$, hence the best-reply is to play $a = H$, too.

Although I drew an analogy between hierarchical archival access in the present formalism and the familiar hierarchical type spaces in the Harsanyi framework, this analogy is merely suggestive. The hierarchies in the present formalism do not represent introspective strategic reasoning or "stepping into the opponent's shoes". Players' epistemology continues to be "flat", in the spirit of self-confirming equilibrium: they approach the problem from a "naively statistical" angle and extrapolate a belief from partial archival data.

5 The False Causation Effect

In simultaneous-move games, each player's action is objectively independent of his opponent's action conditional on his information. In the examples of Sections 4, players' subjective beliefs obeyed this conditional-independence property. As a result, while their beliefs departed from rational expectations, they preserved the rationality property that neither player thinks he can influence his opponent. In contrast, in the example of Section 2, players' subjective action-consequence mapping implicitly involved the wrong belief that their own action affects their opponent's. In this section I present a simple example that makes this error explicit.

A Prisoner's Dilemma example

The two players face 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' archival access. Each player i's news access is fixed: $n_i = \{\mathbf{n}_i, \mathbf{r}_i\}$. Therefore, it can be omitted from the description of the state of the world, which can be written as $\omega = (r_1, r_2, a_1, a_2)$. The distribution over players' archival access is as follows. With probability $1 - \alpha$, $r_1 = r_2 = \{\{\mathbf{r}_1, \mathbf{r}_2, \mathbf{a}_1, \mathbf{a}_2\}\}$. With the remaining probability α, $r_1 = r_2 = \{\{\mathbf{a}_1, \mathbf{a}_2\}\}$. The interpretation is as follows. The variables a_1, a_2 are publicly observed, and therefore the players always know their joint steady-state distribution. In contrast, archival data about the variables r_1, r_2 is "classified"; it gets "declassified" with probability $1 - \alpha$.

This game has an equilibrium in which each player i chooses D *if and only if* $r_i = \{\{\mathbf{r}_1, \mathbf{r}_2, \mathbf{a}_1, \mathbf{a}_2\}\}$. The reasoning is roughly the same as in Section 2. When $r_i = \{\{\mathbf{r}_1, \mathbf{r}_2, \mathbf{a}_1, \mathbf{a}_2\}\}$, player i has rational expectations and therefore plays the dominant action. When $r_i = \{\{\mathbf{a}_1, \mathbf{a}_2\}\}$, $p_{r_i}(a_j \mid r_i, a_i) = p(a_j \mid a_i)$. The putative equilibrium implies $p(a_j = C \mid a_i = C) = p(a_j = D \mid a_i = D) = 1$. Player i's conditional belief effectively interprets this perfect correlation between a_i and a_j causally, as if player j will always mimic player i's choice. Consequently, C is a subjective best-reply for player i.

As in Section 2, the source of the player's error in this example is that he lacks archival access to a variable that acts as a *confounder* between his action and another payoff-relevant variable. Given the underlying assumption that players' understanding of the situation is entirely given by the partial statistical data at their disposal, their subjective beliefs neglect this confounding effect. As a result, the subjective conditional distribution of a_j effectively misinterprets the observed correlation between a_i and a_j as a causal effect of the former on the latter.

One could argue that the source of the player's error is not his first-stage omission, but the fact that when incorporating the second-stage conditional distribution into his choice of action, he imposes a spurious causal interpretation on it. However, computing a conditional distribution and plugging it into an expected-utility calculation is exactly what the player would do in the standard model of Bayesian games. The difference is that when player i forms a conditional belief in the standard model, he conditions on every variable that affects a_i. In contrast, the first stage of the present belief-formation procedure may lead to a violation of this principle.[11]

6 Related Literature

The literature contains a number of equilibrium concepts for static games that are based on the idea that players' equilibrium beliefs are based on partial feedback regarding the equilibrium distribution. It is helpful to define these proposals by the way they formalize partial feedback and the belief-extrapolation rule they assume. The crucial novelty of this paper compared with existing approaches is that it includes an explicit model of the randomness of players' feedback, as well as their uncertainty and limited feedback regarding other players' feedback. This sub-section is designed to clarify the

[11] For an alternative, decision-theoretic approach to this type of "magical thinking", see Daley and Sadowski (2017).

immediate context of this modeling innovation.

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 exhibits "parsimony". 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 are induced by each action. Players ignore sampling error and believe that the sample associated with each action is perfectly representative of its true conditional distribution over outcomes. Salant and Cherry (2019) extend this basic idea to more general estimation procedures. Osborne and Rubinstein (2003) study a variant of this concept, in which each player's feedback consists of a sample drawn from the *unconditional* distribution over the opponent's actions. In Esponda (2008), the feedback sample is infinite but *selective*. For example, in a bilateral trade example, it is the distribution of outcomes conditional on trade. Players' extrapolated belief reflects unawareness of the sample's selectiveness.

Jehiel (2005) and Jehiel and Koessler (2008) present a formalism that is closest in spirit to the present paper, in the sense that a player's feedback limitation is a *personal characteristic* rather than part of the definition of the solution concept. Under this approach, each player best-replies to a coarse representation of the true equilibrium distribution. Specifically, the player partitions the set of possible contingencies 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.

To see how analogy-based expectations and its variants fit into the present framework, consider the following example. Let (x^1, x^2) be the external state. The realization $r_1 = \{\{\mathbf{x}^1, \mathbf{x}^2\}, \{\mathbf{x}^2, \mathbf{a}_2\}\}$ indicates that player 1 learns the distribution of the external state as the joint distribution of player 2's action and a coarse description of the external state (given by the component x^2). The maximum-entropy extension of these marginals is $p_{r_1}(x^1, x^2, a_2) =$

$p(x^1, x^2)p(a_2 \mid x^2)$. This is what the notion of *analogy-based expectations* in static games (Jehiel and Koessler (2008)) would prescribe when x^2 is defined as the analogy class of the external state.[12] When we omit x^2 from the model such that $r_1 = \{\{\mathbf{x}^1\}, \{\mathbf{a}_2\}\}$, we obtain $p_{r_1}(x^1, a_2) = p(x^1)p(a_2)$ - i.e., a belief that a_2 is independent of the external state. This is an instance of "fully cursed" beliefs (Eyster and Rabin (2005)).

The literature also contains a number of models (e.g., Piccione and Rubinstein (2003), Eyster and Piccione (2013), Eliaz et al. (2018)) in which agents' beliefs can be described in terms of archival access that consists of a single subset of variables that omits some relevant variables. In macroeconomics, such beliefs appear in so-called "restricted perceptions equilibrium" (see Woodford (2013)).

In Esponda and Pouzo (2016), players do not extrapolate a belief from limited feedback. Instead, they arrive at the game with a misspecified prior model, and fit this model to their feedback. Esponda and Pouzo 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.

Eyster and Rabin (2005) adopt a different interpretation of distorted equilibrium beliefs. In "fully cursed" equilibrium, player i wrongly believes that the distribution over a_j is a measurable function of player i's signal. In "partially cursed" equilibrium, a player's belief is a convex combination between the correct and fully cursed beliefs. Eyster and Rabin regard this belief dis-

[12] Mailath and Samuelson (2018) assume a similar belief-formation model in the context of an information-aggregation problem.

tortion 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 (2017) provides a partial feedback-based justification for partially cursed beliefs.

7 Conclusion

Previous equilibrium models with non-rational expectations treated belief distortions as an aspect of the solution concept or as a permanent fixture of players. The present formalism enriches the scope of this literature by viewing players' limited learning feedback as an aspect of their type. It describes limited feedback in terms of access to "archival data" about relevant variables. Because players' archival access may itself be a random variable, this language enables us to capture new and realistic kinds of "information about information". It also enables us to explore the economic implications of correlation between this novel aspect of players' types and more conventional aspects.

A natural next step is to examine dynamic strategic interactions, where a move by one player can determine another player's archival access at a later decision node. Eliaz et al. (2021) is a step in this direction: a cheap-talk model in which the sender controls the receiver's news information as well as his access to archival data regarding the joint distribution of messages and states of Nature. We explore the implications of this novel feature on the sender's ability to persuade the receiver.

References

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

[2] Braess, D. (1968). Uber ein Paradoxon aus der Verkehrsplanung. Unternehmensforschung, 12(1), 258-268.

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

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

[5] Daley, B. and P. Sadowski (2017), Magical Thinking: A Representation Result, Theoretical Economics 12.2, 909-956.

[6] Eliaz, K., R. Spiegler and H. Thysen (2021), Strategic Interpretations, Journal of Economic Theory 192: 105192.

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

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

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

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

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

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

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

[14] Eyster, E. and G. Weizsacker (2016), Correlation Neglect in Portfolio Choice: Lab Evidence, mimeo.

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

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

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

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

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

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

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

[22] Milgrom, P. and N. Stokey (1982), Information, trade and common knowledge, Journal of economic theory 26, 17-27.

[23] 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.

[24] 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.

[25] Osborne, M. (2004), An Introduction to Game Theory, Oxford University Press.

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

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

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

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

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

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

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

[33] Spiegler, R. (2020), Can Agents with Causal Misperceptions be Systematically Fooled? Journal of European Economic Association 18, 583-617.

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