

ANALOGY-BASED EXPECTATION EQUILIBRIUM AND RELATED CONCEPTS: THEORY, APPLICATIONS, AND BEYOND *

Philippe Jehiel[†]

24th September 2021

Abstract

A unified definition of analogy-based expectation equilibrium (ABEE) for strategic environments involving multiple stages and private information is presented. Various alternative interpretations of the concept are proposed as well as a discussion of how to use ABEE in practice. A variety of applications including two new ones related to speculative trading and personnel economics is reviewed. A discussion of a number of alternative equilibrium concepts follows emphasizing the links and differences with ABEE. Finally, a discussion of possible next steps in particular related to the endogeneization of analogy partitions is proposed.

Keywords: analogy-based expectations, bounded rationality, limited learning.

*This draft was prepared for the (virtual) world congress of the Econometric Society, Milan 2020. I wish to thank Kai Barron, Milo Bianchi, David Ettinger, Antonio Guarino, Steffen Huck, Frédéric Koessler, Konrad Mierendorff, Dov Samet, Larry Samuelson, Juni Singh, Giacomo Weber for fruitful collaborations on related topics, as well as Ken Binmore, Olivier Compte, Ignacio Esponda, Drew Fudenberg, Simon Gleyze, Laure Goursat, Johannes Hörner, Ariel Rubinstein, Rani Spiegler and seminar participants at various webinars for useful conversations or comments. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 742816).

[†]PSE, 48 boulevard Jourdan, 75014 Paris, France and University College London ; jehiel@enpc.fr

1 Introduction

Standard solution concepts in game theory (and economic theory) typically assume that players (or economic agents) hold rational expectations about the consequences of their actions and about the behaviors of others. While it seems implausible that economic agents would be able to reach this state of knowledge by pure introspection,¹ an appealing approach to the rational expectation assumption relies on learning: As evidence is accumulated, economic agents would gradually correct their mistakes, and eventually behave as the standard concepts assume.

Yet, in complex environments, expectations concern so many different situations that it is unlikely that economic agents would be able to collect enough evidence to correct every conceivable mistake.² I would argue that in such environments, economic agents (similarly as statisticians or analysts) are more likely to form their view of the world based on aggregate statistics that summarize coarsely the behaviors of others.³ It is then of importance to understand how economic agents forming their expectations based on such aggregate statistics would interact in strategic settings, and how the induced behaviors differ from those predicted in the rational case.

In the past years, I have developed a solution concept, the Analogy-Based Expectation Equilibrium (ABEE), aimed at modelling how economic agents exposed to statistics that aggregate (behavioral) data across different situations would behave in steady state. In essentially all applications of ABEE so far, how data are aggregated for players is taken exogenously. That is, one characteristic of a player is how behavioral data are bundled for him, as summarized by his analogy partition. The analogy partition together with

¹Even putting plausibility aside, the pure introspective approach to the rational expectation assumption raises conceptual difficulties, as discussed in Binmore (1987).

²Strictly speaking, a given economic environment is never encountered twice in exactly the same conditions, thereby making the pooling of some evidence/data coming from different sources somehow necessary. Savage (1954) makes the related observation that states as envisioned by decisions makers must be simplifications of the objective states viewed as complete descriptions of the world (which are necessarily unique). This leads him to adopt a subjective formulation of the state space that he does not relate to the objective states. One distinctive feature of the approach developed below consists in linking the subjective representation to objective characteristics.

³Such statistics would be obtained from the pooling of behavioral data corresponding to past play in different situations, and they may either be passed by third parties to new inexperienced players or else correspond to simplified processing of the raw data about past play constructed by the players themselves to facilitate learning. See further discussion below.

the utility function and the private information provides a complete description of the characteristics of the player. Interactions are next described as usual, making it precise who moves when and what the available actions are.

I have considered various strategic environments, covering normal form games played in parallel as well as extensive form games of complete information (Jehiel, 2005), static games of incomplete information (Jehiel and Koessler, 2008), and also some classes of multi-stage games with asymmetric information (Ettinger and Jehiel, 2010). Below, I consider all these settings at once, assuming that the private information is defined at the start of the interaction and that the past actions of players are publicly observable within the game. In such settings, the analogy partition of a player is a collection of subsets of contingencies referred to as analogy classes where a contingency specifies the player j under consideration, the profile of players' types τ , and the history h of past actions.⁴ When two contingencies $(j, \tau h)$ and $(j', \tau' h')$ belong to the same analogy class, I require that the action space of player j at h is the same as the action space of player j' at h' so that there is a well defined action space in each analogy class (this is to make the pooling of the behavior of j at τh and j' at $\tau' h'$ operational in a simple way).⁵ In an ABEE, every player i is assumed to have learned, in each of his analogy classes α_i , the aggregate behavior, referred to as the analogy-based expectation $\beta(\alpha_i)$. Given his private information and his utility, player i plays a best-response to the belief that player j in contingency $(j, \tau h)$ behaves according to $\beta(\alpha_i)$ where α_i is the analogy class to which $(j, \tau h)$ is assigned. While I discuss several motivations for this notion of best-response, my favorite one is based on simplicity: in ABEE, players consider the simplest view of the world that is consistent with their (partial) knowledge, as given by β . That is, player i assumes the behavior in every contingency $(j, \tau h)$ of any of his analogy classes α_i matches the aggregate behavior $\beta(\alpha_i)$ in α_i he is informed about, and he best-responds to this view of the world.

In steady-state, I also require that for each player i and each possible analogy class α_i , $\beta(\alpha_i)$ correctly represents the aggregate behavior in the analogy class.⁶ The correctness of

⁴As the word partition indicates, different analogy classes have an empty intersection and their union is the set of all contingencies.

⁵A possible elaboration of this would view the player as subjectively identifying actions across different contingencies. I will briefly discuss this later on.

⁶To cover analogy classes that would not be reached in equilibrium, I consider trembling-hand per-

the analogy-based expectations is viewed as the outcome of a collective learning process in which before, playing the game, an individual subject assigned to the role of player i receives feedback concerning past behaviors in the various analogy classes α_i corresponding to his type. He chooses his strategy as a best-response to the induced belief about others' behaviors as constructed from the empirical distributions, thereby generating new behavioral data. When a steady-state is reached, the play must be that of an ABEE.

In the rest of the paper, I present more precisely the ABEE concept, develop further the interpretations of it, and discuss applications with a focus on two new ones related to speculative trading and personnel economics respectively, and a highlight on how the ABEE approach can explain puzzles (such as the use of absolute auctions or the emergence of bubbles) or new phenomena (such as deception) that would be hard to capture in the standard rationality benchmark. I also provide some guidance on how to choose analogy partitions in the various applications. In particular, I distinguish three different motivations for the choice of analogy partitions. In the first one, the feedback received by players is viewed as being exogenously coarse, thereby viewing the choice of analogy partitions as being dictated by the kind of coarse statistics about past play that is passed to new players. In the second one, the analogy partitions are viewed as representing what players find more salient in the various contingencies, and this can be used to formalize various psychological biases such as correlation neglect, selection neglect or the fundamental attribution error, among others. In the third one, the analogy partitions are viewed as simplifying learning devices used by players to interpret more effectively the raw data available to them.⁷ After the exposition of ABEE and its applications, I describe how ABEE relate to a number of other concepts including the Bayes Nash equilibrium with subjective prior, the self-confirming equilibrium, the cursed equilibrium, the Bayesian network equilibrium, the Berk-Nash equilibrium, and other equilibrium approaches with some bounded rationality element. Finally, I suggest what I regard as possible next steps for the ABEE research agenda, including possible routes to endogenize the analogy partitions.

turbations in the vein of Selten (1975). I note that such a refinement is not used in most applications of ABEE.

⁷That categorizations are viewed as facilitating learning has, of course, many precedents (see, in particular, Anderson (1991)).

2 The analogy-based expectation equilibrium

2.1 Strategic setup

I consider multi-stage games with observed (past) actions and possibly incomplete information. This allows me to cover simultaneously the setups studied in Jehiel (2005), Jehiel and Koessler (2008) and Ettinger and Jehiel (2010).⁸ Each player $i \in I = \{1, \dots, n\}$ can be one of finitely many types $\tau_i \in T_i$. Player i knows his own type τ_i but not the types τ_j of the other players $j \neq i$. Nature initially selects the profile of types $\tau = (\tau_i)_{i=1}^n$ according to a probability distribution $p(\tau)$, which is commonly known to all players (see below for a discussion of applications in which $p(\cdot)$ is not known to the players).

The game involves finitely many stages. Nature as well as the players may be called to play at various stages of the game. The sequence of possible moves is the same whatever the profile of types τ .

A history refers at the various stages of the game to the string of earlier actions chosen by the various players (as well as nature if applicable), but it does not include the initial selection of types τ . It will be denoted by h , and the set of h will be denoted H . A complete history that includes the initial choice by nature of τ as well as h will be denoted τh . The set of τh will be denoted by TH .

Player i plays at the same subset H_i of histories H whatever his type τ_i . The action space of player i at history $h \in H_i$ is common to all types τ_i , and it is denoted $A_i(h)$. Observe that h may belong to H_i for several i , thereby allowing me to cover the case of simultaneous moves. The set of terminal histories z is denoted by Z , and each player is endowed with a von Neumann-Morgenstern utility function u_i defined over lotteries over Z that may vary with $\tau = (\tau_i, \tau_{-i})$. Specifically, $u_i(z; \tau_i, \tau_{-i})$ denotes player i 's payoff if the terminal node z is reached and the profile of types is $\tau = (\tau_i, \tau_{-i})$.⁹

⁸I am considering games of perfect recall and perfect monitoring. Allowing for imperfect recall and/or imperfect monitoring raises no conceptual difficulties but requires heavier notation (multi-self approaches would also be required to deal with imperfect recall, see Jehiel (2021) for an illustration of ABEE in a context with imperfect recall).

⁹When intermediate payoffs are obtained in the course of the interaction, one should think of $u_i(z; \tau_i, \tau_{-i})$ as aggregating all these intermediate payoffs possibly through the standard exponential discounting formula or any other formula that sounds appropriate.

Notion of type

The non-standard aspect of the strategic environment lies in the definition of types τ_i . A type τ_i of player i consists of two components $\tau_i = (\theta_i, c_i)$ where $\theta_i \in \Theta_i$ is the standard type of player i that in particular can affect players' preferences (payoff-irrelevant signals -such as signals about θ_{-i} - can be part of θ_i as well), and $c_i \in C_i$ that specifies the aggregate statistics (about the behaviors of the various players) player i has access to. I will refer to c_i as the cognitive type or analogy partition of player i .

As seems natural to assume, the cognitive types of players do not affect players' preferences over the various terminal nodes. That is, for each terminal node $z \in Z$, it holds that $u_i(z; \tau_i, \tau_{-i}) = u_i(z; \tau'_i, \tau'_{-i})$ whenever $\tau_i = (\theta_i, c_i)$ and $\tau'_i = (\theta_i, c'_i)$ have the same standard type θ_i , and $\tau_{-i} = (\theta_{-i}, c_{-i})$ and $\tau'_{-i} = (\theta_{-i}, c'_{-i})$ have the same standard type θ_{-i} .¹⁰

Cognitive type

The cognitive type or analogy partition c_i is defined as follows. It is a partition of the set of contingencies described as:

$$\{(j, \tau h) \text{ with } j \in I, \tau \in T, \text{ and } h \in H_j\}.$$

A typical element of partition c_i will be denoted α_i and referred as an analogy class of player i . I impose that if $(j, \tau h)$ and $(j', \tau' h')$ belong to the same analogy class α_i , the action space of player j at h is the same as the action space of player j' at h' . That is, $A_j(h) = A_{j'}(h')$. One can then speak of the common action space $A(\alpha_i)$ in the analogy class α_i .

The idea behind the cognitive type c_i of player i is that this player would only be informed about the aggregate distribution of behavior in each of the analogy classes α_i in c_i , and he would not be informed of more detailed aspects of others' strategies. In order to speak of aggregate behavior, one need to be able to identify actions across different contingencies $(j, \tau h)$ in α_i , and the common action space requirement is a simple way to

¹⁰An alternative commonly used representation of the standard type θ_i involves an underlying state space Ω endowed with some probability distribution, and θ_i is the information set of player i informing this player in which subset of Ω the true state $\omega \in \Omega$ lies. With this representation, the underlying utility would depend on ω (possibly with some aspects of ω being payoff-irrelevant).

make this operational.

Observe that unlike in Jehiel (2005), I am allowing player i to also include player i 's own contingencies in his analogy partitions (as $(j, \tau h) \in \alpha_i$ with $j = i$ is permitted). This should not be confused with the fact that within the interaction I am assuming player i has perfect recall (he remembers his past actions as well as his strategy). Allowing player i to include contingencies of the form $(i, \tau h)$ in an analogy class α_i will allow me to capture applications in which the feedback received by player i about past play aggregates the (past) behaviors of subjects in the roles of players i and $j \neq i$ in various contingencies (see the learning interpretation below and the anonymous analogy partition for an illustration).

2.2 Solution concept

Analogy-based expectations

An analogy-based expectation for player i with analogy partition c_i specifies for each analogy class α_i of c_i , a probability distribution over the action space $A(\alpha_i)$. I will let $\beta(\alpha_i) \in \Delta A(\alpha_i)$ refer to this distribution, and $\beta_{c_i} = (\beta(\alpha_i))_{\alpha_i \in c_i}$ refer to the analogy-based expectation of player i with analogy partition c_i . I will also let $\beta^i = (\beta_{c_i})_{c_i \in C_i}$ refer to the profile of analogy-based expectations for all possible cognitive types of player i , and $\beta = (\beta^i)_{i \in I}$ refer to a profile of β^i for all players $i \in I$.

Strategy

A (behavioral) strategy of player i is denoted σ^i . Following standard approaches, it is a mapping that assigns for every type τ_i of player i and every history $h \in H_i$ at which player i must move, a probability distribution over $A_i(h)$. I let $\sigma_{\tau_i}(h) \in \Delta A_i(h)$ refer to this distribution, and refer to $\sigma_{\tau_i} = (\sigma_{\tau_i}(h))_{h \in H_i}$ as the strategy of type τ_i . A strategy of player i can then be viewed as $\sigma^i = (\sigma_{\tau_i})_{\tau_i \in T_i}$. The set of σ^i will be denoted Σ^i , $\sigma = (\sigma^i)_{i \in I}$ will denote the strategy profile, and Σ the set of σ .

Best-response

Knowing the aggregate behavior $\beta(\alpha_i)$ in α_i (and no more detailed information within α_i), it seems natural that player i would expect player j at τh with $h \in H_j$ to behave according to $\beta(\alpha_i)$ whenever $(j, \tau h) \in \alpha_i$. Formally, this leads player i with cognitive type

c_i to adopt the following representation of the strategy of player j . For every $j \neq i$ and every $\tau h \in TH$ with $h \in H_j$:

$$\widehat{\sigma}_j[c_i; \beta](\tau h) = \beta(\alpha_i) \text{ whenever } (j, \tau h) \in \alpha_i,$$

for each analogy class α_i of c_i .¹¹

In multi-stage games of incomplete information, one classic difficulty is that some histories may not be reached, making the inference about the types of other players at such histories and the ensuing continuation best-response calculation non-trivially defined. I adopt the approach of (extensive-form) trembling-hand proposed by Selten (1975) to get around this difficulty.¹² Specifically, in a first step, I assume, for every player i with type τ_i that every action $a_i \in A_i(h)$ of player i at every history $h \in H_i$ must be played with a probability no smaller than $\varepsilon(a_i, h, \tau_i)$ assumed to be strictly positive.¹³ I refer to the corresponding best-response as an ε -perturbed best-response where $\varepsilon = (\varepsilon(a_i, h, \tau_i))_{a_i, h, \tau_i}$. In a second step I consider sequences of $\varepsilon = (\varepsilon(a_i, h, \tau_i))_{a_i, h, \tau_i}$ such that all $\varepsilon(a_i, h, \tau_i)$ converge to 0.

Formally, ε -perturbed best-response is defined as follows:

Definition 1 σ^i is an ε -perturbed best-response to β^i if and only if for every $\tau_i = (\theta_i, c_i)$, σ_{τ_i} is a best-response to $\widehat{\sigma}_{-i}[c_i; \beta]$ where the optimization must respect the lower bounds imposed by ε .

In Definition 1, I have not made explicit whether the optimization is made at an ex ante stage or at each history where player i must move. Given that all histories are reached whenever all $\varepsilon(a_j, h, \tau_j)$ are strictly positive, and that I am considering games of perfect recall, these two notions of best-responses are equivalent here.¹⁴

¹¹Observe that I consider only players $j \neq i$ because player i need not form a theory about himself (he remembers his strategy (perfect recall) and he does not interact with other subjects in the role of player i).

¹²I should add that the trembling-hand refinement is not needed in most applications of ABEE, to the extent that in most applications all analogy classes are reached in equilibrium.

¹³As in Selten (1975), I assume there is no correlation across the trembles at different nodes.

¹⁴A possible alternative definition of best-response would view player i at different nodes as different players, and each such player could form a belief about the other players in the role of player i using the subjective representation derived from β . In contexts in which a given subject assigned to the role of player i is viewed as choosing an entire plan of actions, the chosen definition is more appropriate.

Consistency

A key equilibrium requirement is the consistency of analogy-based expectations. Specifically, I require that, for every possible analogy class α_i , the analogy-based expectations $\beta(\alpha_i)$ should correctly represent the aggregate behavior in α_i , as implied by the strategy profile σ (and how nature selects types). To define this formally, I let $P^\sigma(\tau h)$ denote the probability that the full history τh is reached when nature selects τ according to $p(\cdot)$ and players behave according to σ . When strategies are fully mixed (i.e., all actions are played with strictly positive probability at every history), $P^\sigma(\tau h) > 0$ for all τh allowing me to unambiguously define consistency as follows.

Definition 2 *Assume that σ is totally mixed. The profile of analogy-based expectations β is consistent with σ if and only if for all possible analogy classes α_i of any player i , it holds that:*

$$\beta(\alpha_i) = \frac{\sum_{(j,\tau h) \in \alpha_i} P^\sigma(\tau h) \sigma_{\tau_j}(h)}{\sum_{(j,\tau h) \in \alpha_i} P^\sigma(\tau h)}. \quad (1)$$

To understand the consistency condition (1), think of a third party (such as a statistician) who would collect information from many different plays about the frequencies of actions played in α_i . Assume also that players keep playing according to σ in all interactions. Such a third party would end up having frequencies that should be a weighted average of the frequencies observed in the various $(j, \tau h) \in \alpha_i$. That is, it should be a weighted average of $\sigma_{\tau_j}(h)$ for the various $(j, \tau h) \in \alpha_i$. Moreover, the weight attached to a specific contingency $(j, \tau h)$ in α_i should be proportional to the frequency of visit of this specific $(j, \tau h)$ as such a frequency indicates how often the behavior in contingency $(j, \tau h)$ will contribute to the observed aggregate frequency profile. The weighting of $\sigma_{\tau_j}(h)$ by $P^\sigma(\tau h)$ is the mathematical formulation of this (given that the frequency of $(j, \tau h)$ in α_i should be proportional to $P^\sigma(\tau h)$). A consequence of this is that if one contingency $(j, \tau h)$ in α_i is much more likely to be visited than another contingency $(j', \tau' h')$ in α_i , $\sigma_{\tau_j}(h)$ will contribute much more than $\sigma_{\tau_{j'}}(h')$ to $\beta(\alpha_i)$, and player i will be induced to believe the behavior of $\tau_{j'}$ at h' is in fact influenced a lot by the behavior of τ_j at h . This extrapolation feature from more frequent contingencies to less frequent contingencies is a key aspect of the analogical expectation formation.

Analogy-based expectation equilibrium

Putting together the best-response and consistency requirements yields the following definitions.

Definition 3 *A strategy profile σ is an ε -perturbed analogy based expectation equilibrium if there exists a profile of $\beta = (\beta^i)_{i \in I}$ such that for every player $i \in I$,*

- 1. σ^i is an ε -perturbed best-response to β^i and*
- 2. β is consistent with σ .*

Definition 4 *A strategy profile σ is an analogy-based expectation equilibrium if there exists a sequence of strategy profiles σ_ε such that for every ε , σ_ε is an ε -perturbed analogy based expectation equilibrium, and σ_ε converges to σ as ε converges to 0.*

2.3 Preliminary properties and learning interpretation

Link to trembling-hand equilibrium

In the special case in which every player i can only have one cognitive type that puts every $(j, \tau h)$ in a singleton analogy class, then it is readily verified that consistency implies that players should have correct expectations, and thus analogy-based expectation equilibria coincide with trembling-hand equilibria.¹⁵

Proposition 1 *When every player uses the finest analogy partition, analogy-based expectation equilibria coincide with extensive-form trembling-hand equilibria.*

Existence

Allowing for arbitrary cognitive types, the existence of analogy-based expectation equilibria is guaranteed in finite environments required to have finitely many types, finitely many actions and finitely many stages.

Proposition 2 *Whatever the environment assumed to be finite, there always exists an analogy-based expectation equilibrium.*

¹⁵It is the extensive-form version of the trembling-hand equilibrium, since trembles were introduced separately at each node.

Proof. The strategy of proof is very similar to that for the existence of trembling-hand equilibria or sequential equilibria (Kreps and Wilson, 1982) in finite environments. For fixed ε , the existence of an ε -perturbed analogy-based expectation equilibrium σ_ε is obtained using Kakutani's fixed point theorem (the consistency mapping is continuous in the admissible range of strategy profiles and the ε -perturbed best-response correspondence is upper-hemicontinuous, thereby allowing for the use of Kakutani's fixed point theorem). Then the compactness of the strategy space in finite environments ensures the existence of accumulation points of such σ_ε when ε goes to 0, and any such accumulation point is an analogy-based expectation equilibrium. \blacklozenge

Learning interpretation

I think of ABEE as the outcome of a learning process. Specifically, think of the game as described above as being played by many different teams of subjects in the role of the various players. Each subject would just play once (similarly as in the recurring game environments considered by Jackson and Kalai (1997)).¹⁶ When a new subject in the role of player i comes in, think of his cognitive type c_i as informing him of the aggregate behavior from past play in each of the analogy classes α_i in c_i . That is, when in a previous play the history τh was reached and j played a_j , this action is added to the α_i dataset where $(j, \tau h) \in \alpha_i$. A player with cognitive type c_i is informed of the frequencies of actions in the various analogy classes α_i of c_i . Such a player then best-responds to the frequencies in the various α_i as in Definition 1, that is, assuming that the strategy of j at τh matches the aggregate frequency in $\alpha_i(j, \tau h)$ where $\alpha_i(j, \tau h)$ is the analogy class in c_i to which $(j, \tau h)$ belongs. His behavior generates new data that are used by future generations.

Such procedures generate dynamic systems that are parameterized by the initial play (and possibly whether and to which extent more remote observations are thrown away from the datasets). The analogy-based expectation equilibria are the steady states of such dynamic systems, since in steady-state, the long-run frequencies observed in each analogy class would have to give rise to consistent analogy-based expectations as given by

¹⁶I could consider that subjects play several times but importantly I have in mind that the aggregate statistics as opposed to the own experience would be the main driver for the choice of strategy.

condition (1) of Definition 2.^{17,18} A complementary study not covered in this paper but discussed at the end as a possible direction for future research is when we should expect such dynamic systems to converge to a steady state.

2.4 Preliminary discussion

Interpretations of the coarse analogy partitions

I suggest here several alternative interpretations of the coarse analogy partitions that I divide in three broad families. In the first one, following the learning interpretation just provided, a player may receive feedback for each of his analogy classes without having any control on the type of feedback he receives. This would fit with situations in which depending on his environment (cultural or otherwise) an agent is exposed to (possibly agent-specific) aggregate statistics. With this interpretation, it is natural to view the analogy partition of a player as part of the description of the environment in the same way as his information partition. I will refer to this as the **exogenously coarse feedback interpretation**.

The second family differs from the first one in that the coarseness of the analogy partition of a player is viewed as a subjective characteristic of this player. Specifically, some aspects of the data describing how players previously behaved in the various contingencies may be more salient than others, and what appears as more salient may be agent-specific. The clustering into analogy classes may then formalize what features players with different types find more salient when collecting data from past interactions (in this interpretation, a player with cognitive type c_i would only recognize from a past contingency $(j, \tau h)$ the analogy class $\alpha_i(j, \tau h)$ to which it belongs and $\alpha_i(j, \tau h)$ can then be interpreted as the salient part of contingency $(j, \tau h)$ perceived by such a player). I will refer to this as the **salience-based interpretation**, and it will be used to model various biases as identified

¹⁷The steady state formulation is reminiscent of the self-confirming equilibrium, which will be discussed later on.

¹⁸The learning dynamics just described has some similarities with the fictitious play dynamics first introduced by Brown (1951) in the context of normal form games with the following essential differences. First, I am considering more complex games here and the frequencies are aggregated over different $(j, \tau h)$ which is the key distinctive feature of the present approach as opposed to the more traditional approach in (fictitious-play) learning. Second, I am considering populations of different teams who would just play once whereas the classic fictitious play dynamics considers that the same players keep playing again and again.

by psychologists.¹⁹

Finally, in the third family, players can have access to all the characteristics of the raw data describing the observed past behaviors in the various contingencies, but players are viewed as simplifying what they try to learn so as to facilitate learning. The fewer the analogy classes, the less a player has to learn about his environment, and presumably the less demanding is the requirement that analogy-based expectations are consistent. I will refer to this as the **simplified learning interpretation**. While I find such a view appealing in light of some basic insights drawn from statistics (in particular in relation to the bias-variance trade-off), articulating a full-fledged theory formalizing it will require further understanding of how players envision this simplification exercise. I will touch upon this toward the end of the paper.

Simplicity of the subjective representations given the feedback

An important aspect of ABEE is that in the face of the aggregate statistics $\beta(\alpha_i)$ in the various analogy classes α_i of c_i , player i best-responds assuming that τ_j in τh behaves according to $\beta(\alpha_i(j, \tau h))$ where $\alpha_i(j, \tau h)$ is the analogy class in c_i to which $(j, \tau h)$ belongs.

The main argument for this assumption is simplicity. When one only knows the aggregate behavior $\beta(\alpha_i)$ in α_i , the simplest theory compatible with this is that in every contingency $(j, \tau h)$ in α_i , the corresponding behavior of player j at τh matches the aggregate behavior $\beta(\alpha_i)$. To make the link to simplicity formal, I can make use of the complexity theory as envisioned in the automaton literature that relates the complexity of a strategy to the number of states its algorithmic (or automaton) representation requires (see Neyman (1985) or Rubinstein (1986)). Indeed, in the ABEE modeling, player i need just one state to describe the strategy in α_i whereas any alternative theory compatible with $\beta(\alpha_i)$ would require more states. In addition, checking whether the considered subjective theory is compatible with β is trivial for the subjective theory considered in ABEE, but it would necessitate non-trivial calculations for any alternative theory, thereby providing another complexity-based argument in favor of the ABEE modeling.

¹⁹In more specific contexts, Bordalo et al. (2012) and Köszegi and Szeidl (2013) propose to structure salience on the basis of the degree of difference of the states as measured by some distance (absolute or relative). Such features can be used to put structure on the analogy partitions in such contexts. More generally, Gardenfors (2000) discuss more extensively how to think of categorization in relation to similarity (see also Rubinstein (1988) for a related view in the context of choices over lotteries).

There are alternative possible rationales that would work in some cases. One of them is related to the principle of insufficient reason. This principle has been formalized in the statistical physics literature by proposing a selection device based on the maximum entropy criterion (see Jaynes (1957)). When the choice of theory (here subjective strategies attributed to others) does not affect the perceived probability of reaching the various elements in α_i , entropy would be maximized by requiring that behaviors are uniformly the same across the various contingencies $(j, \tau h)$ in α_i , as in the ABEE modeling. However, in extensive form games and for general analogy partitions, maximizing the entropy may result in quite complex computations, thereby making this rationale less compelling as a general principle.²⁰

What players need to know in an ABEE

In many (if not most) game theory applications, it is argued that the structure of the game - preferences, information structure and game tree - is common knowledge among players. While the common knowledge assumption is not required even for classic notions of equilibrium when one thinks of equilibria as the outcome of learning processes,²¹ it is of interest to address explicitly what players need to know about their strategic environments

²⁰To give a concrete illustration of this, assume a player must decide at node N_1 whether to Pass or Take and if he Passes (thereby reaching node N_2), this player must again decide whether to Pass or Take. Assume the two nodes N_1 and N_2 of this player form a single analogy class for other players, and let $p \in (\frac{1}{2}, 1)$ denote the aggregate probability of Pass in this analogy class.

Letting p_i be the conjectured probability of Pass at node N_i , consistency of (p_1, p_2) with p boils down to

$$p = \frac{p_1 + p_1 p_2}{1 + p_1}$$

which given that $p > \frac{1}{2}$ requires that

$$p_2 = \frac{p - (1 - p)p_1}{p_1} \tag{2}$$

is strictly positive.

The resulting probability of reaching Take after node N_1 is $1 - p_1$, Take after node N_2 is $p_1(1 - p_2)$ and Pass after node N_2 is $p_1 p_2$.

The resulting entropy viewed as a function of p_1 after replacing p_2 by expression (2) $Ent(p_1)$ has a derivative (w.r.t. p_1) equal to

$$-p \ln p - (1 - p) \ln(1 - p) > 0$$

at $p_1 = p$. Thus, $Ent(p_1)$ is not maximized at $p_1 = p$, where $p_1 = p$ would correspond to the belief in the ABEE. In order to maximize the entropy, p_1 should be set above p so as to better equalize the probabilities that all three alternatives are reached.

²¹Such a claim is illustrated for example by Kalai and Lehrer (1993). In their learning model, one gets convergence to Nash equilibrium while players need not know anything about the preference of others.

in the presence of types with coarse analogy partitions. In particular, do subjects need to know the distribution of analogy partitions of other players as well as other aspects of the strategic environment?

Given the interpretation of the consistency condition in terms of learning, I would argue that players need to know very little about their environment. In particular, a given player i with cognitive type c_i need not have any knowledge about the implication of the cognitive types of others, nor does he need to know the preferences or the implication of the private information for others. Player i need only know his own preference structure u_i , the game tree as well as the distribution of the type profiles τ . He should also know the aggregate behavior in each of his analogy classes α_i in c_i . But, importantly, as already highlighted, I regard the latter knowledge as the result of learning coming from the processing of behavioral data observed in previous interactions. As a result, this knowledge does not require any understanding (or knowledge) about how other players reason.

It should be mentioned that, in the above setting, the knowledge of u_i as well as of the distribution $p(\cdot)$ of types is required so as to allow player i to compute the best-response as specified in Definition 1. Viewing Nature as a player (as I do in some applications) would allow player i to not know $p(\cdot)$, and player i would then form a subjective perception about $p(\cdot)$ in the same way as he currently forms a subjective perception about σ_{-i} (using a coarse analogy partition). Such uses of ABEE will be discussed later on.

Is there a universal type space?

The notion of type in ABEE is non-standard given the cognitive dimension c_i . From a theory viewpoint, one may be interested in exploring whether one can define a notion of universal type space that would include the cognitive dimension in the same way as it is done for classic Bayesian games (see Harsanyi (1967) and Mertens and Zamir (1985)). It turns out that one cannot in general define a universal type space (at least if one sticks to the representation of types as proposed above). The reason for this makes use of the classic result of Cantor who has established the impossibility of constructing bijections between the set of subsets of a set and the set itself. Indeed, given that a cognitive type can be viewed as specifying partitions of the cognitive types of others (as well as other elements such as the regular type θ and the node h), it is readily verified using Cantor's

theorem that it is not possible to define a universal type space. This is different from the usual approach to types in Bayesian games in that an information type in the classic case is a partition not over the type space of others but over an appropriately defined state space so that one can avoid the self-referential aspect of the partitioning in the classic case but not in the present context.²²

I should add that I do not regard this theoretical observation as problematic from an applied viewpoint. Every strategic environment corresponds to a given set of types and the cognitive types correspond to different ways of pooling data from past observations. The relevant set of types is then dictated by the application one has in mind. To give a concrete illustration of this, in many applications, it would seem natural that the cognitive type of a player is never observed even ex post, in which case analogy partitions can be reduced to partitions of $(j, \theta h)$ where θ refers to the profile of regular types. In this case, the self-referential problem related to Cantor's theorem would disappear and the set of all types can be defined unambiguously.²³

Similarity in actions

In the above framework, I have assumed that when two contingencies $(j, \tau h)$ and $(j', \tau' h')$ belong to the same analogy class α_i of player i , the action space of player j at h should be the same as the action space of player j' at h' , i.e., $A_j(h) = A_{j'}(h')$. I motivated this as a way to introduce more simply the notion of aggregate distribution of actions in the analogy class α_i . This view of common objective action space seems particularly legitimate when considering the exogenously coarse feedback interpretation of the approach in which the labels of actions would typically chosen by the third parties in charge of passing the aggregate statistics.²⁴ However, with the other proposed interpretations of analogy

²²Heifetz and Samet (1998) offer a related even if a bit different perspective on a related issue. They show that one cannot define a universal knowledge space (unlike for probabilistic belief spaces as considered by Harsanyi); and they attribute this impossibility to a lack of continuity in the knowledge operator. A formal investigation in the line of Heifetz and Samet for the present context should be the subject of further research.

²³Even if c_i is not observed in applications, sometimes some aspects of θ_i can be correlated with c_i . In such applications, analogy partitions of $(j, \theta h)$ can have an effect similar to those obtained when a partition of $(j, \tau h)$ is considered.

²⁴Even with this interpretation, one could envision situations in which the action spaces in different contingencies of α_i are not the same and yet player i is informed of the aggregate distribution of behaviors in α_i . In this case, I would suggest that, for $(j, \tau h) \in \alpha_i$, player i expects player j at τh to behave according to the aggregate distribution $\beta(\alpha_i)$ conditional on the action of j being in $A_j(h)$.

partitions, one could envision more general situations in which players could subjectively consider various identifications of actions across different contingencies. This would lead to enrich the cognitive type of player i so as to include a specification for this player of which action of j at h is identified with which action of j' at h' .²⁵

Link to Nash Equilibrium

A question of theoretical interest is whether the ABEE is a novel concept or whether after an appropriate modification of the game it can be viewed as a familiar equilibrium such as the Nash equilibrium. To shed light on this question, I consider the following simple setup. There are two possible stage games G_1 and G_2 chosen with equal probability by Nature. These games are symmetric, involve two players $i = 1, 2$ and two actions a and b with player 1's payoff as shown in the following tables (where player 2's payoffs can be inferred by symmetry):

G_1	a	b	G_2	a	b
	a	x	z	a	t
	b	y	t	b	z

Assume that players observe the stage game they are in, and that they use the coarsest analogy partition (putting the two games together for each of the players). Playing a in game G_1 and b in game G_2 is a symmetric ABEE whenever $x + z \geq y + t$. This follows because for the analogy-based expectation β to be consistent with this strategy profile, it should be that β assigns equal probability to a and b (remember that both games are equally likely), and a (resp. b) is the best-response to such β in G_1 (resp. G_2) whenever $x + z \geq y + t$.

Observe that this ABEE is not a Nash equilibrium whenever $y > x$.²⁶ It is not either the Nash equilibrium of the game in which players would not be informed whether the stage game is G_1 or G_2 , as with such an information structure, players would have to

²⁵It would then lead to the weaker requirement that the cardinality of the actions spaces across different contingencies are the same, except if one is willing to consider the possibility that different actions in the same contingency are considered as similar to one another by the player, as in the work of Arad and Rubinstein (2019) in which case no restriction would be imposed on the cardinality.

²⁶In game G_1 , players choose a and not b which is the best-response to a in this game because they (wrongly) expect their opponent to be playing according to β , thereby leading them to (subjectively) find a optimal.

behave in the same way in G_1 and G_2 , which is not so in the ABEE just constructed.^{27,28}

Of course, if one allows oneself to change the payoffs x, y, z, t to alternative values x', y', z', t' , one can view the ABEE with payoffs x, y, z, t as a Nash Equilibrium with payoffs x', y', z', t' , for example by letting $x' = 1$ and $y' = z' = t' = 0$. But, this is a tautological statement as any behavior can be rationalized in such a way.²⁹

A more interesting inquiry is when there are several ABEE, whether a single change of payoffs and information structure is enough to identify the set of all Nash equilibria in the modified setting with the set of all ABEE in the original setting. The answer to this question is negative, as illustrated in the above class of symmetric setup with coarse analogy partition when $x > t > y = z = 0$. In this case, the symmetric pure ABEE are such that either a is played in both games, or that b is played in both games, or that a is played in game G_1 and b is played in game G_2 . But, note that b in game G_1 and a in game G_2 is not an ABEE, the reason being that if it were so, the analogy-based expectation β would have by consistency to assign equal probability to a and b , and the best-response to this would be a and not b in game G_1 (given that $x > t$). Now, consider a modified setting in which the payoffs can be changed (I will refer to G'_1 and G'_2 as the corresponding stage games) and the players may or may not be informed of the stage game. Clearly, given that in at least one ABEE, the behaviors are not the same in G_1 and G_2 , players need to know in which game they are. But, then no matter how x', y', z', t' are chosen, it would

²⁷In general, the ABEE cannot be related either to the Bayes Nash equilibrium in which players have subjective prior about the information of their opponent (here whether or not they know the stage game). I do not develop this here, as the Bayes Nash equilibrium with subjective prior will be discussed later.

²⁸A somewhat dual investigation one could consider is whether the cognitive types c_i alone could capture the various forms of asymmetric information as usually modeled by θ_i (I thank Jannes Hörner for suggesting this question). The answer is negative, thereby illustrating that the two dimensions in the type play a different role. To illustrate this, consider the same two game scenario and assume that players do not know the game G_1 or G_2 they are in. It is readily verified (exploiting the symmetry) that the mixed strategy which assigns equal probability to a and b in each game and for each player would be a Bayes Nash equilibrium. Given that such play would be the same across all contingencies, consistency would imply that players must have the correct expectation about their opponent's behavior no matter how analogy partitions are defined. Yet, if players know the game they are in, they would strictly prefer a in game G_1 whenever $x + z > y + t$ with such a correct expectation, thereby establishing that the proposed strategy profile cannot be an ABEE under complete information no matter how the cognitive types are defined.

²⁹For a solution concept like ABEE that preserves symmetry (G_2 is obtained from G_1 by exchanging the roles of a and b), it is enough to simply consider changes of x, y, z, t to x', y', z', t' that themselves preserve symmetry.

not be possible to rule out the behaviors b in G'_1 and a in G'_2 as a Nash equilibrium when the other three combinations would be Nash equilibria.³⁰

Consistency in extensive form games

A key aspect of ABEE lies in the consistency condition, as defined in (1). In particular, the endogenous weighting arising in extensive form games can give rise to new theoretical properties that would not arise had the weighting been exogenously fixed. To illustrate this, consider the following simple two-player, two-stage game with complete information. Player 1 moves first and chooses an action L or R . After observing player 1's action (L or R), player 2 chooses an action l or r . Assume that it is strictly optimal for player 2 to choose r after L and to choose l after R . Assume that $(2, L)$ and $(2, R)$ form a single analogy class α for player 1. The interaction is depicted in the following figure.




FIGURE 1

In the standard rationality case, L or R would be perceived by player 1 to be optimal in the Subgame Perfect Nash Equilibrium. Under the conditions shown below, the coarse analogy partition of player 1 forces him to randomize in ABEE.

³⁰The modifications just considered assume that the payoffs \tilde{u}^j in stage game G_j can only depend on the actions a^j chosen in that game. If one instead allows for the possibility that the payoff of a player in game G_j depends also on the actions in game G_{-j} more can be done. Specifically, consider the normal form game with four players denoted (i, j) with $i = 1, 2$ and $j = 1, 2$ in which the payoff of player (i, j) is

$$\tilde{u}_i^j(a_i^j, a_{-i}^j, a_{-i}^{-j}) = p_j u_i^j(a_i^j, a_{-i}^j) + p_{-j} u_i^j(a_i^j, a_{-i}^{-j})$$

where a_i^j denotes the action of player (i, j) identified with the action of player i in Game G_j in the original setting. If p_j is the probability of G_j and $u_i^j(a_i, a_{-i})$ is the utility of player i in Game G_j when the actions are (a_i, a_{-i}) in the original setting, it is readily verified that the Nash equilibria of the modified normal form game coincide with the ABEE of the original setting.

The reason why such a transformation can be done in a way that works for all ABEE is that in the case of static games as considered in Jehiel and Koessler (2008), the probabilities of the various elements of an analogy partition can be determined independently of the equilibrium. Such a transformation would not in general work in multi-stage settings in which such probabilities would be endogenously determined by the equilibrium (as reflected by the terms $P^\sigma(\tau h)$ in the consistency condition (1)).

Claim. If $u_1(Rl) < u_1(Ll)$ and $u_1(Lr) < u_1(Rr)$ there is no ABEE employing pure strategy: Player 1 must be mixing in ABEE.

Proof. Assume in ABEE, player 1 always chooses L . Consistency implies that $\beta(\alpha) = r$, but then player 1's best-response to β is R , since $u_1(Lr) < u_1(Rr)$. So player 1 cannot be choosing L with probability 1 in ABEE. Similarly, if player 1 were to always choose R , then $\beta(\alpha) = l$ by consistency, which would lead player 1 to choose L by the best-response application. This shows that there is no ABEE in pure strategy. The ABEE (which exists by Proposition 2) must be in mixed strategy. ♦

Observe that if the weighting in the consistency formula had been fixed independently of the strategy of player 1 (for example, putting equal weight on the behavior of player 2 after both L and R), then (at least) one of the two actions L or R would have been perceived by player 1 to be optimal, and having player 1 play this action with probability 1 would have been part of an equilibrium. It is the endogeneity of the weighting that forces player 1 in ABEE to be randomizing.³¹ To go further in understanding the shape of the randomization in this setting, consider the following payoff matrix that indicates the payoffs of players 1 and 2 for each action profile:

$$\begin{array}{cc} & l & r \\ L & (x, 0) & (z, 1) \\ R & (y, 1) & (t, 0) \end{array}$$

with $x > y$ and $t > z$ so that the conditions of the claim are satisfied.

It is readily verified that there is a unique ABEE in which action L is played with probability $\frac{x-y}{t-z+x-y}$ and action R is played with probability $\frac{t-z}{t-z+x-y}$ so that the aggregate distribution of behavior of player 2 is $\beta = \frac{x-y}{t-z+x-y}r \oplus \frac{t-z}{t-z+x-y}l$ -a lottery assigning probability $\frac{x-y}{t-z+x-y}$ to r (which is observed when player 1 chooses L) and probability $\frac{t-z}{t-z+x-y}$ to l (which is observed when player 1 chooses R). Note that the mixing of player 1 in the ABEE is chosen so that the induced aggregate behavior of player 2 as given by the consistency condition makes player 1 feel he is indifferent between choosing L or R . As

³¹More generally, and beyond the simple environment of Figure 1, the weighting of the contingencies in the analogy-based expectation cannot be viewed as exogenous, as soon as some strategic decision is to be made prior to a contingency in the analogy class.

a result, in the ABEE, player 1 chooses more often action L when x and z are increased, which stands in contrast with what comes out in the Nash equilibrium of the normal form game in which player 1 would choose L and R with an equal probability 0.5 irrespective of x , y , z and t (the insensitivity of the strategy to one own's payoffs is sometimes regarded as a counterintuitive prediction of the Nash equilibrium in mixed strategies).

From another perspective, if we have in mind that subjects observe fully the profile of past plays in similar interactions, in the proposed ABEE, the observation would result in the (L, r) profile with frequency $\frac{x-y}{t-z+x-y}$ and in the (R, l) profile with frequency $\frac{t-z}{t-z+x-y}$. The proposed ABEE setting would then model subjects in the role of player 1 who are subject to correlation neglect, missing the correlation between the actions of players 1 and 2 and trying only to construct the marginal distribution of behaviors of subjects assigned to the role of player 2 from the data they observe. I will discuss further the use of the ABEE setting to model correlation neglect as well as other psychological biases.

ABEE in the lab

While this paper almost exclusively emphasizes the theoretical properties as well as the various applications of ABEE, I regard the experimental validation of the approach as a necessary step. The experimental validation is still in its infancy, and my hope is that more tests will appear in the future. In Huck et al. (2011), we have provided the simplest experimental test of ABEE. Subjects either in the role of Column player or in the role of Row player were faced with two possible stage games that they played overall the same proportion of time. The feedback that subjects in a given role (Column or Row) received concerned the past play of subjects assigned to the other role, and it could either be coarse in the sense that the behaviors over the two stage games were aggregated or fine in the sense that this feedback was cleanly distinguished according to the game. There were additional treatments in which the feedback consisted of raw data of past play with no statistical treatments. The results in the basic treatments were quite in line with the ABEE predictions and illustrative of the exogenously coarse feedback interpretation of the concept. The treatments with raw data illustrated that when data are made difficult to process, in some instances, subjects focus on simplified statistics such as those resulting from aggregating data from different situations (as in ABEE) in line with the interpretation in terms of simplifying learning device.

This study is encouraging that the ABEE may be useful to describe the behavior of agents exposed to coarse (or hard to process) feedback. But, what this study is not addressing (because it would require richer environments) is when subjects are exposed to raw data and there are several possible ways of aggregating data, what bundling device subjects use and when (see however Grimm and Mengel (2012) for some hints related to that). It is also not touching on the effect of having multiple stages and private information.³²

2.5 On the use of analogy partitions

Clearly, several possible analogy partitions can be considered. This is somehow similar to the observation that different information partitions can be considered in the classic Bayesian setting. As for the modeling of information, I would say the application and context one has in mind should inform the modeler what kinds of analogy partitions are more relevant. I suggest some hints below.

2.5.1 Some classes of analogy partitions

With the exogenously coarse feedback interpretation in mind, some analogy partitions emerge more prominently depending on what seems more natural to assume about what is observed or not from past similar interactions. In particular, the following analogy partitions seem natural in a number of applications.

Anonymous analogy partitions

Consider situations in which all players play at the same time and share the same action space. Sometimes, when observing past behaviors from previous interactions, it is not known whether player i or j was behind the action. This is particularly relevant when the feedback one obtains about past play is anonymous. In such a case, it would seem natural to consider anonymous analogy partitions in which, for any h and τ , all permutations of $(i, \tau h)$ with respect to the set of players would belong to the same analogy

³²I have been involved in a few other experiments in relation to ABEE, which include Ettinger and Jehiel (2021) as a test of the theory of deception developed in Ettinger and Jehiel (2010), and Barron et al. (2020) as a test of Jehiel (2018) (see also Jehiel and Singh (2021) for a test of Valuation equilibrium, which is closely related to ABEE, as explained below).

class. I have considered such an analogy partition in the context of asymmetric First-Price Auctions (Jehiel, 2011).

History-independent analogy partition

Consider multi-stage interactions with complete information in which the action space of players is the same at all stages. Sometimes, when being told the actions from past interactions, it is not known after which history (sequence of actions) the actions were chosen. In such cases, it is natural to consider history-independent analogy partitions defined so that, for each player, all histories at which the player must move are bundled into one analogy class. The personnel economics application discussed later on will illustrate this.

Payoff-relevant analogy partition

In a number of contexts, it is difficult to have access to what previous agents engaged in similar interactions knew at the time they made their decisions. In such cases, it is natural to consider the payoff-relevant analogy partition defined so that all $(j, \tau h)$ corresponding to the same payoff-relevant information but possibly different information at the time of the interaction are bundled in the same analogy class. I have considered such an analogy partition in the context of social learning (Guarino and Jehiel, 2013) and of auction mechanisms (Jehiel and Mierendorff, 2021).

In situations in which a natural distance $d(\cdot, \cdot)$ can be defined on the (standard) state space Θ , it is not uncommon that behaviors observed in nearby states would be pooled together when constructing aggregate summary statistics, thereby leading to the following family of analogy partitions.

Neighborhood analogy partition

Consider a set of representative elements θ^k in Θ , and assign every θ in Θ to the closest category k defined such that $d(\theta, \theta^k) = \min_r d(\theta, \theta^r)$. A neighborhood analogy partition for player i is such that $(j, \tau h)$ and $(j, \tau' h)$ belong to the same analogy class whenever θ and θ' as appearing in τ and τ' belong to the same category k . This is related to the

clustering idea familiar in machine learning, and I will consider this idea again toward the end of the paper.

A different idea of bundling by closeness is for each θ to consider a neighborhood of θ , call it $N(\theta)$, in Θ , and view player i as expecting the behavior in θ to match the aggregate behavior in $N(\theta)$ (this can be related to the k nearest neighbors algorithm familiar in machine learning, and Steiner and Stewart (2008) can be viewed as considering such a bundling device in the context of coordination games).³³ I would argue that this latter form of bundling by closeness is better suited to model situations in which the player himself being informed of the state θ would decide how to bundle the data whereas the neighborhood analogy partition as previously defined describes better situations in which the aggregate statistics are made available by third parties without knowing which state will be drawn.

2.5.2 Link to cognitive biases

While the above examples of analogy partitions were motivated by the kind of information about past interactions that is passed to new comers, analogy partitions can also be used to formalize a number of biases as identified by psychologists. Some examples follow. They can be viewed as illustrations of the salience-based interpretation of the analogy partitions.

Extrapolation bias

ABEE can directly be linked to the idea of extrapolation. For example, if players are engaged in two different types of interactions say A and B involving the same action space in both, one captures the idea that behaviors observed in either interaction are extrapolated to both interactions by putting the two interactions in the same analogy class. If one interaction is more frequent than the other, the behavior in that interaction has more weight on the overall expectation, as follows from the consistency requirement.

³³While this would not lead to a partitioning of Θ given that the same state can belong to several such neighborhoods (mathematically, $N(\theta) \cap N(\theta')$ need not be empty when $\theta \neq \theta'$), it should be mentioned that, if one considers that a player i in state θ is not the same player as player i in state θ' , one can represent such situations using the (analogy) partitioning formulation. This simply requires endowing player i in state θ with an analogy partition in which one analogy class bundles the decision nodes of all players j in states $\theta' \in N(\theta)$ (the other analogy classes for this player are irrelevant, since player i 's utility depends only on what happens in the realized state θ).

I have considered such a use of ABEE in environments with Bargaining and Ultimatum interactions, as explained later.

Correlation neglect

Another common bias of interest in behavioral economics is related to correlation neglect (see Ellis and Piccione (2017) for some recent study of correlation misperception or Razin and Levy (2019) for a recent survey on the implications of correlation neglect in political economy, see also the cursed equilibrium of Eyster and Rabin (2005) to be discussed later for a model of correlation neglect in strategic environments). To illustrate it, consider the following simple setting. Two variables, say b and c , are correlated and distributed according to $f(b, c)$. Agents may fail to appreciate the correlation and reason as if these variables were independently distributed. A natural way to model this is to consider that (b, c) would be perceived to be distributed according to the product distribution $\tilde{f}(b, c) = f_b(b)f_c(c)$ where $f_x(x)$ denotes the marginal distribution of $x = b$ or c , as derived from the joint distribution $f(b, c)$.

Such misspecifications can easily be accommodated in the framework of the analogy-based expectation equilibrium. For simplicity, assume that b and c are variables chosen by nature even though in general one or both variables could be decisions made by some players (as is the case, for example, in the cursed equilibrium). Consider the following extensive form game in which Nature is decomposed into three different players referred to as Nature, N_b and N_c . At the root of the interaction Nature chooses ω defined to be (b, c) and drawn according to $f(., .)$ (assumed not be known by the decision maker, see the above comments on the possibility that players would not know the prior distribution of types). Then N_b and N_c move simultaneously. N_b chooses $b = b(\omega)$ and N_c chooses $c = c(\omega)$ where $b(\omega)$ and $c(\omega)$ denote the first and second coordinate of ω , respectively. Other player(s) may then proceed in their interaction based on how (b, c) as chosen by N_b and N_c is distributed.

If a player puts all contingencies (N_b, ω) into one analogy class α_b and all contingencies (N_c, ω) into another analogy class α_c , the perceived distribution of (b, c) will be given by $f_b(b)f_c(c)$ in equilibrium (as follows from the application of consistency), thereby showing how correlation neglect can be captured using the analogy-based expectation equilibrium

framework.

Fundamental Attribution Error (FAE)

FAE as coined by Ross (1977) is "the tendency to believe that what people do reflects who they are" or more generally the tendency to pay insufficient attention to the context in which actions are performed to form beliefs about the kind of person one is facing. In Ettinger and Jehiel (2010), we have proposed to model FAE using the ABEE framework. In an environment with several groups x of agents and several groups $g \in G$ of interactions, if all (g, x) with $g \in G$ are bundled into one analogy class α_x to assess the behavior of agents in group x , one precisely implements the idea that agents reason based on the aggregate attitude of agents in group x and make inferences based on this aggregate knowledge. In other words, this type of ABEE modeling captures the bias consisting in thinking that what people do is only determined by their group x and not by the situation g they are in, which seems very much in line with how social psychologists informally refer to the FAE.

Selection Neglect

Selection bias has been recognized in statistics and econometrics as an important source of mistakes in estimation exercises. When the selected sample that one observes is not representative because say it is determined by endogenous decisions made by agents, a good analyst should be cautious when extrapolating from observed data (see for example Heckman (1979) for an influential analysis on how to correct for selection bias).³⁴ One mistake or bias suggested by psychologists is that not only bad analysts but also (some or even many) human beings themselves may be victims of selection neglect (see, for example, Koehler and Mercer (2009)).

Selection neglect can be captured in the analogy-based expectation equilibrium framework. For concreteness, think of a decision maker who has to choose between $a = 1$ (spending one more year at school) and $a = 0$ (dropping from school). For the sake of

³⁴A classic example is related to the estimation of returns to schooling. If one extrapolates from those who attend extra years of schooling to estimate the benefit of these extra years, one ignores that those attending these extra years are generally not random subjects but subjects who thought these extra years would be beneficial, thereby leading to systematic biases in the estimations.

illustration, let $\theta = (b, c)$ denote a state characterizing the decision problem where the cost c but not the benefit b attached to $a = 1$ is observed by the decision maker. In a standard description of the interaction, nature chooses θ according to distribution $f(\cdot)$, then the agent observes c and chooses $a = 0$ or 1 with payoff consequences given by $u(a; \theta) = b - c$ if $a = 1$ and $u(a; \theta) = 0$ if $a = 0$.

A rational agent would choose $a = 1$ whenever $E(b | c) > c$. By contrast, a selection neglect agent would assess the distribution of b by extrapolating from the distribution of b observed from other agents who chose $a = 1$. To model this mistaken behavior with ABEE, consider the following alternative formulation of the interaction now described as a two-player game. After Nature chooses θ and the agent chooses $a = 1$, player N_b chooses $b = b(\theta)$ where $b(\theta)$ denotes the first coordinate of θ . By putting all the $(N_b, (\theta, a = 1))$ contingencies in a single analogy class so as to assess the distribution of b as chosen by N_b , the agent will in equilibrium form an expectation about b identical to that of a selection neglect agent. That is, only those c for which $a = 1$ is chosen will contribute to the expectation formation, and the agent will choose $a = 1$ whenever $E(b | a = 1) > c$. Interestingly, even though we are in a decision problem, there is an equilibrium construction to the extent that $E(b | a = 1)$ determines the instances c such that $a = 1$ is found to be the optimal decision, and the decision strategy in turn affects $E(b | a = 1)$.³⁵ Jehiel (2018) provides an illustration of such a selection-neglect construction in investment decision contexts.

Other biases

Other biases can be modelled using the analogy-based expectation equilibrium. For example, the bias consisting in attributing to others one own's information can be modeled by considering the private information analogy partition in which $(j, \tau h)$ and $(j, \tau' h)$ are assigned to the same analogy class by player i whenever $\tau_i = \tau'_i$. Indeed, if player i believes that other players share his own information, it is then natural to assume that player i will bundle the behavior data of player j according to player i 's own information, hence the private information analogy partition. Note that in static environments, this corresponds to the fully cursed equilibrium of Eyster and Rabin (2005) that I will discuss

³⁵The dependence of the bias on how c affects the decision a makes this bias distinct from correlation neglect as usually considered, even if at some general level, one can view selection neglect as a special form of correlation neglect (in which the variable b would be viewed as independent of a and c).

later on.³⁶ One may also capture the idea that player i would believe he has superior information by endowing player i with an analogy partition that would be coarser than player i 's information partition (as given by τ_i).

In other contexts, player i may believe that player j ' strategy depends only on some characteristics of the history, for example what is payoff-relevant or the recent history, thereby leading player i to bundle $(j, \tau h)$ according to such characteristics. When the retained characteristic is the payoff-relevance criterion, this may seem related to the Markov Perfect Equilibrium (MPE). Yet, a key difference is that in the ABEE modelling unlike the MPE modelling, the premise that the strategy of j depends only on what is payoff-relevant need not be correct. In a similar vein, the literature on sunspot equilibria assume that some payoff-irrelevant aspects (the sunspots) affect behaviors, but, like for MPE, this assumption is correct in equilibrium, making the sunspot equilibrium closely related to the (publicly) correlated equilibrium rather than the analogy-based expectation equilibrium. It turns out that there are two early illustrations of that kind of ABEE modelling that were developed before the first publication on ABEE appeared. One of them is the moonspot equilibrium discussed by Aziariadis and Guesnerie (1982) toward the end of their paper in which some economic agents think some variables depend on moonspots when in reality they depend on other characteristics. The other one is a multi-stage monopoly problem studied by Piccione and Rubinstein (2003) who consider consumers who only recognize patterns of prices of some fixed length l .³⁷

3 Illustrations

To show possible applications of the analogy-based expectation equilibrium, I start with two illustrations specifically developed for this survey. The first one can be viewed as a stylized study of how traders basing their strategy on coarse statistics may be led to trade in contexts in which the no trade theorem would normally apply. It is an elaboration

³⁶See also Madarasz (2012) for an alternative modelling of this bias that he refers to as information projection.

³⁷While Piccione and Rubinstein' setup can be captured in the ABEE setting by assuming that histories h are categorized according to their last l action profiles, it should be mentioned that the resulting ABEE equilibrium can also be viewed as an equilibrium with imperfect recall of the multi-self sort analyzed by the same authors in an earlier publication (Piccione and Rubinstein, 1997).

of the betting game example discussed in Jehiel and Koessler (2008). The second one considers a personnel economics application in which the coarse statistics the employer has access to leads sophisticated employees to engage into deceptive tactics in probation periods in an attempt to obtain more discretion in the subsequent phase of the labor relationship. It is an elaboration of the monitoring example discussed in Ettinger and Jehiel (2005) (which is the working paper version of Ettinger and Jehiel (2010)). While the former application is concerned with a static Bayesian environment, the latter considers a multi-stage environment with multiple types, thereby providing illustrations of ABEE in strategically diverse environments. I end the Section with a quick presentation of other ABEE applications, as appearing in the literature. A reader willing to see a discussion of the connection of ABEE to other approaches should jump to Section 4.

3.1 On speculative betting

Strategic betting environment

Two players $i = 1, 2$ are engaged in betting games involving four possible scenarios referred to as $k = e, b, c, d$. Players 1 and 2 decide simultaneously to bet or not bet, $a_i \in \{B, D\}$ for $i = 1, 2$. If either player chooses D , the bet is not implemented and each player gets 0. If both players choose B , the bet is implemented and the payoffs of the players depend on the scenario k as well as the realization of some random variable $x = (x^k)_{k=e,b,c,d}$. Specifically, in scenario k with x realization, $u_1(B, B; k, x) = x_1(k) = x^k$ and $u_2(B, B; k, x) = x_2(k) = -x^k$.

Players are informed of the realized x but not of the scenario k . About k , player $i = 1, 2$ can either receive the signal that k lies in $\{e, b\}$ or $\{c, d\}$ -this information structure will be referred to as \bar{I} - or that k lies in $\{e, c\}$ or $\{b, d\}$ -this information structure will be referred to as \underline{I} . A profile of information structures (I_1, I_2) for the two players will be referred to as I .

I will be assuming that each scenario k can arise with the same probability $1/4$, that the information structure of each player is equally likely to be \bar{I} or \underline{I} and that the probability that the two players have the same information structure (i.e., $I_1 = I_2$) is $\gamma \in [\frac{1}{2}, 1]$ where $\gamma = 1/2$ corresponds to the case of independence and $\gamma = 1$ corresponds to the case of perfect correlation in the draws of the information structures of the two players.

Moreover, each x_k for $k = e, b, c, d$ is drawn independently of each other from the same distribution, assumed to be symmetric around 0 with no mass on 0.

In any Bayes Nash Equilibrium of this game, no bet can possibly be implemented in equilibrium. This follows because the interaction is zero-sum and every player can guarantee an ex-ante 0 payoff by always choosing D no matter what the private information is. The formal analysis of such games appears in Sebenius and Geanakoplos (1983). It can be seen as an illustration of the so called no trade theorem (see Milgrom and Stokey (1982) for a general formulation of it).

Consider now the following cognitive environment. Each player $i = 1, 2$ can have only one cognitive type c_i so that I can identify a state ω with (x, k, I) .³⁸ Moreover, the analogy partition considered by player $i = 1, 2$ consists in partitioning the set of $(i, \omega = (x, k, I))$ according to whether $x_i(k) > 0$ or $x_i(k) < 0$.³⁹

Such an analogy partition can be interpreted as considering that players are informed (for past similar plays) about the aggregate probability of bet whenever it was in the ex-post interest of the player to bet (this corresponds to the class including (i, ω) when $x_i(k) > 0$) - call this aggregate probability β^+ -, and the aggregate probability of bet whenever it was not in the ex-post interest of the player to bet (this corresponds to the class including (i, ω) when $x_i(k) < 0$) - call this β^- . Or to use the language of statistics, this can be interpreted as players receiving feedback from past play on the proportion of betting errors of types 1 and 2 (with the notion of errors defined at an ex-post stage when full information on ω is available). Having in mind that such two-type-error statistics are somehow salient when describing the behaviors of previous players, I would like to suggest that this analogy partition is of particular interest in this application.

Analysis

Given β^+ and β^- assumed to be strictly positive, the strategy choice of players is straightforward. Consider to fix ideas player 1, and assume that he receives the signal that k is in $\{e, b\}$ as well as observing x . If $x^e > 0$ and $x^b > 0$, clearly player 1 will choose B . Similarly, if $x^e < 0$ and $x^b < 0$, player 1 will choose D . The non-trivial case is when

³⁸As usual, one can link the ω formulation to the type θ formulation used in Section 2 by identifying θ_i with (x, I_i) .

³⁹Contingencies $(i, \omega = (x, k, I))$ with $x_i(k) = 0$ are irrelevant given that the distributions of x^k have no mass on 0.

$x^e = x^+ > 0$ and $x^b = x^- < 0$ (or vice versa). In this case, best-response to $\beta = (\beta^+, \beta^-)$ requires that player 1 chooses B whenever

$$\beta^- x^+ + \beta^+ x^- > 0 \tag{3}$$

and D otherwise.⁴⁰ This is because in scenario e , player 2 would be getting $-x^+ < 0$ in case the bet is implemented and thus player 1 perceives player 2 to be betting with probability β^- in such a case. Similarly, in scenario b , player 1 perceives player 2 to be betting with probability β^+ . Since the two scenarios e and b are equally likely, playing B is perceived to yield $(\beta^- x^+ + \beta^+ x^-) / (\beta^- + \beta^+)$ in expectation to player 1, hence the choice of strategy.

Despite the simplicity of the betting decision criterion (3), it implicitly hides some non-trivial inference process. While a fully naive player would bet whenever $x^+ + x^- > 0$ given that the two scenarios are a priori equally likely, here player 1 infers from his knowledge of β^+ and β^- that if player 2 chooses to bet, the odds that the scenario is e rather than b is shifted from 1 to $\frac{\beta^-}{\beta^+}$, hence leading to criterion (3).

A key question of interest then is how β^+ compares to β^- in equilibrium, and whether and when the bet is implemented. In general, β^+ and β^- are determined by the consistency conditions which are influenced by the chosen distribution of x_k . For example, assume that each x_k is drawn from a distribution that is either the uniform distribution on $[-1, 1]$ with probability $1 - \mu$ or the distribution putting equal mass $1/2$ on ε and $-\varepsilon$ (with ε being arbitrarily small) with probability μ . Solving for the fixed point in (β^+, β^-) that

⁴⁰I am implicitly assuming that in case of indifference, the player chooses D .

results from the consistency conditions for the two analogy-based expectations yields:⁴¹

$$\begin{aligned}\beta^- &= -1/4 + (1 - \mu^2)/8 + ((1/2 - ((1 - \mu^2)/4)^2 + \mu(1 - \mu))^{1/2} / 2 \\ \beta^+ &= 1/2 + \beta^-\end{aligned}$$

When $\mu = 0$ or 1 , $\beta^- = 0$ and $\beta^+ = 1/2$ so that there is no bet that is implemented in equilibrium. But, when $0 < \mu < 1$, we have $0 < \beta^- < 1/2 < \beta^+ < 1$, and thus some bets are implemented in an ABEE. For example, when $\mu = 1/2$, one gets $\beta^- \approx 0.14$ and $\beta^+ \approx 0.64$ so that a non trivial bet is implemented when player 1 has information structure \bar{I} , player 2 has information structure \underline{I} , scenario $k = b$ occurs, $x^e = x^+$, $x^b = x^-$, $x^c < 0$ or $|x^c/x^b| < \beta^-/\beta^+$.

Moving away from the special distributions just considered, one can show that $\beta^+ \geq \beta^-$ quite generally, for any distribution, and for any ABEE.⁴² This observation implies that when both players have the same information, i.e. when the information structure of both players is the same $I_1 = I_2$, no bet is implemented in the ABEE, as necessarily one of the two players would find D optimal.⁴³ Observe also that the consistency conditions remain the same independently of the correlation γ in the information held by the two players so that keeping all other parameters fixed, the share of bets that are implemented in an ABEE decreases as the correlation γ of the two players' information structures increases. Such qualitative properties match the intuitive idea that in order to induce some speculative trade (here some betting activity), a necessary feature is that agents

⁴¹The consistency conditions can be written as

$$\begin{aligned}\beta^+ &= \frac{3}{2}\mu(1 - \mu) + \frac{\mu^2}{2} + \frac{(1 - \mu)^2}{2}\left(1 + \frac{\beta^-}{2\beta^+}\right) \\ \beta^- &= \frac{\mu(1 - \mu)}{2} + \frac{(1 - \mu)^2}{2}\frac{\beta^-}{2\beta^+}\end{aligned}$$

from which the expressions of β^- and β^+ can be derived.

⁴²This is because if $\beta^+ < \beta^-$, (3) would imply that there is betting when $x^+ \geq x^-$ and given the independence of distributions, this would imply by the consistency condition that $\beta^+ \geq \beta^-$, thereby leading to a contradiction.

⁴³To see this, suppose both players are informed that $k \in \{b, e\}$. If x^k for $k = b, e$ have the same sign, the player who would receive negative payoff from the bet must choose D so that the bet is not implemented. When the signs of x^b and x^e are different, the fact that one of the players chooses not to bet is an immediate consequence of the observation that $\beta^+ > \beta^-$.

should observe different pieces of information. Observe that this same conclusion need not hold in environments with subjective priors (the formal definition of which will be reminded later on), as sometimes proposed to get around the no trade theorem.⁴⁴

Turning to the specific distribution of x^k as considered above, and considering the magnitude of β^- as a measure of the betting activity, one can see that β^- is minimized when $\mu = 0$ or 1 , and maximized for an interior value of μ . Interpreting the values ε and $-\varepsilon$ as insignificant ones, the possibility of insignificant $x^k = -\varepsilon$ when associated with $x^{k'} > 0$ in an information structure in which k and k' are not distinguished leads to an increase of β^- , which in turn allows for some betting activity. While very stylized, this illustrative example is suggestive that the possibility of insignificant payoffs may in some cases have important overall economic consequences (here in terms of betting activity) through the indirect effect of modifying the overall expectations of agents (as a result of the analogical reasoning).

3.2 On manipulative behavior in probation periods: A personnel economics application

Strategic environment

Consider an interaction between an employer or Principal (referred to as P) and an employee or Agent (referred to as A) consisting of two phases 1 and 2. Phase 1 is the probation phase in which P (the principal) watches the behavior of A (the agent) when there is no contractual arrangement between them. At the end of phase 1, the principal decides on the contractual arrangement which consists in either giving the agent discretion (action D for Delegation) -this will be the same as the no contractual arrangement prevailing in phase 1-, or exerting tight monitoring (action C for Control). Phase 1 lasts n_1 periods and phase 2 lasts n_2 periods. I have in mind that $n_2 > n_1$. Some of the analysis will be presented for the case in which n_1 and n_2/n_1 tend to ∞ .

In each period of both phases, the agent has to decide on a (productive) effort level $e \in [0, 1]$ where $1 - e$ can be thought of as the non-productive effort. When there is no

⁴⁴Indeed when the signs of the relevant x^k are not the same, if each player thinks the state which is favorable to him is sufficiently more likely, betting can occur in a Bayes Nash equilibrium with such subjective priors. See Liang (2020) for a recent study of learning allowing for speculative trade along such lines in complex enough environments.

contractual arrangement as in phase 1 or when the principal has chosen D in phase 2, the stage game payoffs in each period are:

$$\begin{aligned} u_P(e) &= \pi \cdot e \\ u_A(e) &= w + b_{PB} \cdot (1 - e) + b \cdot e \end{aligned}$$

for P and A , respectively. In these expressions, w can be interpreted as the per period wage received by the agent, e as the productive effort of the agent, $1 - e$ as the non-productive effort, $b \cdot e$ as the bonus received by the agent in proportion of his productive effort, $b_{PB} \cdot (1 - e)$ as the private benefit received by the agent in proportion of his non-productive effort and πe as the payoff (assumed to be proportional to the productive effort of the agent) obtained by the principal. I will be assuming that $b_{PB} > b$.

When contract C prevails (in phase 2), the stage game payoffs of the Principal and the Agent are respectively:

$$\begin{aligned} u_P(e; C) &= \pi - c \\ u_A(e; C) &= w + b \cdot e \end{aligned}$$

where c can be interpreted as the cost of control incurred by the principal. In the C regime, whatever productive effort e the agent (initially) chooses, control leads him to execute the task (as resulting from productive effort $e = 1$) while losing the private benefit attached to the non-productive effort with a bonus proportional to his initial productive effort. In the C regime, the principal safely guarantees π irrespective of e , but this comes at the cost of control c .

Both the Principal and the Agent are aggregating strings of stage game payoffs according to the standard exponential discounting formula. I will consider two types of agents: the patient ones who have a discount factor close to 1 and the impatient ones who have a discount factor close to 0 (and accordingly, only care about their current stage game payoff). The ex ante probability that the agent is patient is $\mu_0 \in (0, 1)$. The Principal on the other hand will be assumed to be patient with a discount factor close to 1. In line with the general class of games considered above, I assume that whether the Agent is Patient or Impatient is not observed by the Principal, and that the profile of effort decisions in

phase 1 is observed by the Principal at the end of the probation phase.

If the Principal and the Agent were rational, the analysis would be as follows. Given that $b_{PB} > b$, the only way for the Principal to incentivize the Agent to exert productive effort is to have the C contract. It follows that at the end of phase 1, the Principal would choose C no matter what happens in phase 1. As a consequence, the Agent irrespective of whether he is patient or not would choose $e = 0$ in all periods of phase 1 (since he would see no impact on his phase 1 behavior on the C/D choice of contractual arrangement). Moreover, in all periods of phase 2, the Agent would choose $e = 1$ (resp $e = 0$) if C (resp. D) is chosen by the Principal.

I will now consider a different cognitive environment in which the Principal forms his expectations based on the aggregate effort attitudes of each type of Agents, not distinguishing those as a function of history. Or in the ABEE language, I will assume that the Principal uses two analogy classes that distinguish behaviors of the Agent according to whether he is patient or not but not according to the history of play (including which contractual arrangement C or D is chosen by the Principal at the end of phase 1). I will further assume that the Agent whether patient or not is rational.⁴⁵

Analysis

The statistics on which the Principal bases her reasoning are β^p , the aggregate effort distribution of patient Agents, and β^i , the aggregate effort distribution of impatient Agents. Note that I have in mind that these statistics concerning other similar interactions have been passed to the Principal, not that the Principal constructs herself these statistics from her own past experience. Such coarse statistics in turn lead the Principal to reason as if she were facing an adverse selection problem in which there are two types of Agents with different effort attitudes (as reflected by β^p and β^i). The effort choices observed in phase 1 (the probation phase) are then used by the Principal to update the likelihood that the Agent is either one of the two types, which is used to determine which of contract C or D looks subjectively best to the Principal. From the perspective of the Agent, the behavior of the Principal dictates the best effort choice to make in phase 1. The nature of the interaction looks then like a signalling game (on Agent' side) with some

⁴⁵Impatient but rational Agents can equivalently be viewed as patient but coarse Agents (who would use a single analogy class to assess the C/D decision of the Principal so that they would not see the effect of their phase 1 behavior on the Principal's choice and accordingly behave as if they were myopic).

screening taking place at the end of the probation phase (on Principal' side). Importantly, the misperception of the Principal is a consequence of the coarse statistics she has access to, which do not allow her to understand the causal effect of the choice of contract on the effort attitude of the Agent.

To analyze this environment, observe first that the behavior of impatient Agents is straightforward. It is $e = 0$ in all periods of phase 1, and $e = 1$ (resp. $e = 0$) in all periods of phase 2 if the Principal chooses contract C (resp. D). Concerning patient Agents, their behavior is the same as that of impatient Agents in phase 2 for each possible contract C or D , but their behavior in phase 1 has to be determined.

Clearly, patient Agents would prefer having the D contract rather than the C contract at the end of phase 1, since such a D contract would allow them in the subsequent periods to obtain $w + b_{PB}$ rather than $w + b$, the stage game payoff they obtain when the contract C is chosen (remember that $b_{PB} > b$).

I will construct an ABEE in which a patient Agent chooses his effort decisions $e = 0$ or 1 in the probation phase so as to obtain the D contract at the end of phase 1, and the Principal chooses contract C after observing the historical effort record of an impatient Agent in the probation phase.⁴⁶ Under such circumstances, consistency would imply that β^i is a distribution assigning probability $\beta^i(1) = \frac{n_2}{n_1+n_2}$ on effort $e = 1$ and probability $\beta^i(0) = \frac{n_1}{n_1+n_2}$ on effort $e = 0$ (given that for impatient Agents, contract C would be chosen on-the-path, thereby leading to the frequency $\frac{n_2}{n_1+n_2}$ of observed effort $e = 1$ for such Agents). Concerning β^p , it would be a distribution on $e \in \{0, 1\}$ with a probability mass on 0 at least equal to $\frac{n_2}{n_1+n_2}$ given that in phase 2, a patient Agent would always be choosing $e = 0$ (after the choice of D contract made by the Principal).

Focusing on the case in which n_2 is large compared to n_1 , impatient Agents are thus perceived to be choosing $e = 1$ with a large probability (close to 1) while patient Agents are perceived to be playing $e = 0$ with a large probability (close to 1). If the Principal

⁴⁶There is another ABEE in which both types of agents behave in the same way and the coarse Principal randomizes between contract C and D so that the induced aggregate effort behavior of agents make the contracts C and D equally attractive to the Principal. Such an equilibrium would be the only ABEE if there were only impatient agents. I note however that in the two type case, such an ABEE is not robust to the introduction of behavioral types who would always choose $e = 1$ in phase 2 (as then patient agents would have an incentive in phase 1 to build a reputation for being this behavioral type in an attempt to obtain the D contract for sure).

thought she was facing an impatient Agent for sure by the end of phase 1, she would choose contract D , expecting to get next $\pi \frac{n_2}{n_1+n_2}$ in each period, rather than contract C that would only give $\pi - c$ to the Principal in each period.

Assuming n_1 is large enough, if a patient Agent were to always choose $e = 1$ in phase 1, he would convince the Principal by the end of the probation period that he is almost surely an impatient Agent, thereby ensuring that the Principal chooses contract D . A patient Agent will not need to exert effort $e = 1$ in all periods of phase 1 though. Given that he myopically likes $e = 0$ better than $e = 1$ (in phase 1), he will choose the minimum number of periods in which he chooses $e = 1$ so as to persuade the Principal to pick the D contract. Call this minimum number n_1^* .⁴⁷ Consistency would then imply that β^p assigns probability $\beta^p(1) = \frac{n_1^*}{n_1+n_2}$ to effort $e = 1$ and the complementary probability to effort $e = 0$.⁴⁸ The above considerations ensure that for $n_1, n_2/n_1$ sufficiently large, one can construct an ABEE with the desired properties.

Comment.

An interesting feature of the proposed ABEE concerns the mistaken belief of the Principal about which type of Agents she is facing at the end of phase 1. When she faces an impatient Agent, she observes n_1 periods with $e = 0$, which makes the Principal very

⁴⁷Given discounting, the patient agent will choose $e = 0$ in the first $n_1 - n_1^*$ periods and $e = 1$ in the last n_1^* periods of the probation phase.

⁴⁸To complete the characterization of the equilibrium value of n_1^* , let μ^* be the threshold belief about the agent being patient such that any lower belief would lead the Principal to choose contract D . Given the aggregate effort attitudes of the two types of agents (as defined by β^p and β^i), the threshold belief μ^* is defined by the indifference condition:

$$\pi(\mu^* \beta^p(1) + (1 - \mu^*) \beta^i(1)) = \pi - c.$$

Given that $\beta^p(1) = \frac{n_1^*}{n_1+n_2}$ and $\beta^i(1) = \frac{n_2}{n_1+n_2}$ would be close to 0 and 1 respectively as $\frac{n_2}{n_1}$ grows large, it is readily verified that μ^* would get close to $\frac{c}{\pi}$ in this limit.

The equilibrium number n_1^* will be the minimum integer k such that

$$\frac{\mu_0}{1 - \mu_0} \frac{(\beta^p(1))^k (1 - \beta^p(1))^{n_1-k}}{(\beta^i(1))^k (1 - \beta^i(1))^{n_1-k}} < \frac{\mu^*}{1 - \mu^*}$$

where the left-hand side of the inequality is simply the posterior likelihood that the agent is patient relative to being impatient after observing k periods with $e = 1$ and $n_1 - k$ periods with $e = 0$ in phase 1. The definition of n_1^* ensures that the patient agent makes the minimum number of effort $e = 1$ decisions so as to trigger the choice D of contract. Assuming n_2/n_1 and n_1 are large, $\beta^p(1) \simeq 0$, $\beta^i(1) \simeq 1$, and it is readily verified that n_1^* should be close to $n_1/2$ irrespective of μ_0 .

confident she is facing a patient Agent (who is perceived to be exerting effort $e = 0$ with very high probability as n_1 and n_2/n_1 tend to infinity). By contrast, when she faces a patient Agent, her belief at the end of phase 1 is that she is facing an impatient Agent (viewed as overall exerting effort $e = 1$ with large probability) with enough probability (so that contract D looks best to her). Moreover, it is precisely by behaving in a way other than that of an impatient Agent in phase 1 that a patient Agent manages to increase the belief in the Principal's mind that he is an impatient Agent overall exerting effort $e = 1$ with a large enough probability. These features illustrate how belief manipulation in ABEE differ from that in traditional reputation/signalling models in which the belief (in others' mind) about being another type necessarily decreases over time in expectation (it is a submartingale). It is an illustration of how ABEE can account for deception and manipulation in a way that is not feasible in the traditional paradigm with full rationality.

Discussion

This simple personnel economics model can be enriched in many directions. I now consider several possible variants and indicate their likely effects while leaving for future research the detailed analysis.

I first consider adding a share of rational Principals that can be viewed as experienced Principals, while still assuming that coarse Principals -viewed as novice Principals- know only the aggregate effort behavior of patient and impatient Agents (across all periods and all possible matches, i.e. whether an Agent is matched with a rational or a coarse Principal). Clearly, rational Principals would choose contract C at the end of the probation period no matter what the observed behaviors are in phase 1, as this is the only way to incentivize the Agent to exert effort and choose $e = 1$ in phase 2. When the share of rational Principals is not too large, one can support an ABEE very similar to the one shown above in the absence of rational Principals. As before, in this ABEE, rational patient Agents adjust their phase 1 behavior so as to obtain contract D at the end of the probation period when they face a coarse Principal. The effect of the rational Principals is to increase the aggregate probability that $e = 1$ for patient Agents as represented by $\beta^p(1)$ given that when patient Agents face rational Principals, patient Agents choose $e = 1$ (and not $e = 0$) in phase 2. This increase of $\beta^p(1)$ in turn forces patient

Agents to choose $e = 1$ more often in phase 1 in an attempt to persuade coarse Principals that it is preferable to go for the D contract.⁴⁹ Thus, the presence of rational Principals induces two costs for patient Agents. On the one hand, they make the manipulation less effective as there is a chance for patient Agents that they face rational Principals. On the other hand, they make the manipulation more costly for patient Agents as more effort periods are now required in phase 1 to persuade coarse Principals to choose the D contract. It should also be mentioned that coarse principals benefit from the presence of rational principals as these induce more effort from patient Agents in phase 1. When the share of rational principals is too large, such an ABEE cannot be sustained anymore, and patient Agents turn to be behaving like impatient Agents. Interestingly, when the share of rational Principals becomes very large, coarse Principals would then believe that all Agents exert $e = 1$ with a large probability and accordingly would go for contract D at the end of phase 1. In this case, the presence of rational Principals would be detrimental to coarse Principals. These various effects illustrate the potential positive or negative externalities that rational Principals can exert on coarse Principals as well as on patient and impatient Agents.

Another possible variant is to allow the Principal at least in some cases to fire the Agent at the end of phase 1 in which case the Principal would hire a new Agent drawn from the pool of Agents (and decide with this freshly drawn Agent whether to choose the D or C contract). Suppose the firing option is only available with some probability z assumed to be not too big (maybe because firing costs are large in a majority of cases). One can construct an ABEE similar to the one shown in the main analysis. Patient Agents will seek to obtain the D contract by an appropriate behavior in phase 1. Firing whenever available will only be used for impatient Agents as these would be observed to have consistently chosen $e = 0$ in phase 1. Moreover, if the share of patient Agents is large, the Principal would choose the C contract in case of new hire in phase 2, since such Agents would be viewed as exerting little effort and these are more numerous. This in turn would increase $\beta^P(1)$ the aggregate probability that patient Agents choose the $e = 1$ effort (while $\beta^i(1)$ would remain close to 1 as $n_1, n_2/n_1$ tend to infinity), and this would

⁴⁹More precisely, as n_1 and $n_2/n_1 \rightarrow \infty$, rational agents would choose $e = 1$ a proportion $\frac{1+\beta^P(1)}{2}$ of the time in phase 1 (where the $e = 0$ choices would all appear at the start of phase 1 due to discounting).

overall induce the patient Agents to choose more often $e = 1$ in the probation phase (so as to obtain the D contract). If on the other hand the proportion of impatient Agents is large, the principal would choose the D contract in case of new hire. This would lower the value of $\beta^i(1)$ (while not affecting much the value $\beta^p(1)$ which would be close to 0 as $n_1, n_2/n_1$ tend to infinity), and this would lead patient Agents to choose $e = 1$ less often in the probation phase.

There are, of course, many extra variants that could be considered including the possibility that the Principal changes the contract in the course of phase 2 (as a response to unexpected effort observations made then) or that Agents may differ in characteristics other than the impatience making the screening at the end of phase 1 valuable even for rational Principals. I hope the above discussion will trigger further work in such directions.

3.3 Other applications

In this part, I review the various applications of ABEE I am aware of (with an obvious bias toward my own work). They are grouped by themes.

Bargaining and Ultimatum

In Jehiel (2005), I have considered two possible applications of ABEE to bargaining. First, I have considered setups consisting of both bargaining and ultimatum interactions where the former differs from the latter in that after an offer is rejected, a counter-offer can be made. I have considered the case in which to assess the rejection rate of offers, both interactions are bundled into one analogy class, thereby formalizing the idea that acceptance/rejection observations made in bargaining would be extrapolated to ultimatum and vice versa. The first finding is that the standard prediction in which a balanced offer is made in Bargaining and an imbalanced one is made in Ultimatum both being accepted with probability 1 does not emerge as an ABEE. Indeed, if it were so, the belief would be that offers whether balanced or imbalanced are accepted in all interactions due to the observation that both balanced and imbalanced offers would be made in equilibrium (although not in the same interaction), and they would always be accepted. Such a belief would lead the proposer to make an imbalanced offer also in Bargaining, as this would be perceived to be more advantageous, and such an offer would in turn be rejected, invalidating the construction of the putative ABEE. The ABEE must take a different

form. If Bargaining is much more common than Ultimatum, one can establish that in an ABEE, a balanced offer would have to be made in both interactions, as an imbalanced offer would be considered to be rejected with high probability due to the analogy with what happens in Bargaining.⁵⁰ Thus, the analogy between Bargaining and Ultimatum may explain why not so imbalanced offers are made in Ultimatum as observed experimentally, without any reference to fairness considerations as usually assumed.

Second, focusing on single bargaining interactions (say of the ultimatum type in which the rejection of the responder would lead to positive even if welfare-inferior payoffs), I have considered the possibility that the proposer would assess the rejection rate of his offers, not for every offer separately but only according to whether the offer is generous or not (say above or below an exogenously given threshold). In such a cognitive environment, the proposer could in ABEE make offers above what would make the responder indifferent between accepting and rejecting (as the traditional analysis would imply). The key to understanding this insight is that, following the logic of the best-response in ABEE, if the proposer decides to make an offer in an analogy class, he would go for the least generous offers among those given that he would predict the same rejection probability for all such offers. The ensuing analysis of ABEE is then very similar to that obtained when the action space of the proposer is assumed to contain only these discrete offers, thereby providing the desired result.

Cooperation in finite horizon interactions and bubbles

In the finitely repeated prisoner's dilemma, the Subgame Perfect Nash Equilibrium (SPNE) predicts that there should be universal defection at all histories. Similarly, in the centipede game (Rosenthal (1982)), players should Take at all nodes in the SPNE. This follows from the logic of backward induction. In the last period, such behaviors correspond to a dominant strategy. Anticipating such behaviors in all subsequent periods of the game implies that the same behaviors should apply in the current period no matter what the history is. The same logic would lead to the conclusion that there can be no bubbles in finite economies (see Tirole (1982)). Such predictions are however not observed in lab experiments (see Selten and Stoeker (1986), Embrey et al. (2018) or Smith et al.

⁵⁰This approach is distinct from the one pursued in Samuelson (2001) who follows the automaton approach by assuming that a player suffers an extra cost when using a strategy that is not the same across different interactions.

(1988) among others).

In Jehiel (2005), I have considered the ABEE in the centipede game assuming that players put all the nodes of their opponent in a single analogy class. Such an analogy partition can be viewed as modeling a scenario in which players would only be informed of the aggregate Pass rate in past similar interactions or alternatively as modeling a possible bias in which players would wrongly consider the strategy of their opponent to be history-independent. A key observation is that some Pass behavior can arise in an ABEE, and the predictions are that either Take occurs early in the game (as in the SPNE) or very late (and more so as there are more periods) but not in the middle of the game.^{51,52} In Bianchi and Jehiel (2010), the same logic has been extended to provide some possible rationale to the emergence of bubbles in finite economies.⁵³ More recently, using the ABEE tool, Antler (2019) has enriched the basic insight of cooperation in the centipede game to provide a rationale for pyramid scams.

Reputation and deception

In Ettinger and Jehiel (2010), we have used ABEE to model the FAE and develop a theory of deception. We have applied it to study a stylized bargaining problem in which the seller of a house by acknowledging small deficiencies with the heating system would be able to persuade the potential buyer that he is more likely to be honest, making the next (more significant) claim that there is another buyer in line more credible. In Jehiel and Samuelson (2012), the same theme has been pursued to revisit the classic theory of reputation (Fudenberg and Levine 1989) in which a long-run player faces a sequence of short-run players always playing the same stage game, and the long-run player may either be rational or mechanical (in which case he always behaves in the same way). The departure from the classic reputation model is that the short-run players are assumed to only know the aggregate behavior of the rational type over all histories

⁵¹To see this, observe that if Take were to occur in a pure strategy ABEE in the middle of the game, one of the players would only be Passing on the path and thus the expectation of the other player as defined by consistency would be that this player passes always leading him to prefer Passing in the middle of the game rather than Taking.

⁵²The observation that the Take phase is shorter as there are more periods is in line with some recent experimental finding derived in the centipede game (see Krockow et al (2018)), see also Danz et al. (2016) for an experiment playing on the feedback and Garcia et al. (2020) for another recent experiment).

⁵³See Moinas and Pouget (2013) for an experimental setting related to bubbles.

and time periods. That is, the short-run players aggregate for each type of long-run players the behavior across all histories into one analogy class. The obtained ABEE differs from the classic analysis in a number of respects. In particular, while traditional models emphasize that reputation motives arise when there is a value to commitment, the analysis with ABEE suggests that reputation building may be at work even when there is no value to commitment (as in zero-sum games) in an attempt to manipulate the belief of boundedly rational agents. The personnel economics application developed above is another illustration of deception and reputation building with analogical reasoning.

Social learning

In the classic social learning model (Banerjee (1992) and Bikchandani et al. (1992)), agents make binary decisions (A or B) in sequence on the basis of a binary signal that is informative about whether A or B is better as well as the observation of previous agents' decisions. When all signals are equally informative, long-run inefficiencies and herding may occur, as the information publicly conveyed by past choices may be superior to the private information. Smith and Sorensen (2000) have modified the basic setup to allow individual agents to receive private information of different precisions. When the precision can be unbounded, they have shown that in the long-run, there are no inefficiencies. The key driving force behind this asymptotic efficiency result is that when an agent breaks a long series of equal decisions (a local herd), the next agents infer that the deviating agent must have had a very precise information going against the crowd, and they are happy to follow him when their own information is not very precise. Smith and Sorensen (2000) refer to this as the overturning principle.

In Guarino and Jehiel (2013), we have considered such social learning environments, assuming that agents make their inferences based on the payoff-relevant analogy partition. That is, agents would have learnt from past similar social learning games, the proportion of time option A was chosen when it was the better alternative and similarly the proportion of time option B was chosen when it was the better alternative, but not the precise distribution of actions as a function of the history.⁵⁴ The ABEE prediction in this case is that even if the precision of signals can be unbounded, there must be inefficiencies in the

⁵⁴I would argue such (coarse) statistics describe the most salient aspects of past decisions, and as such they are more likely to be accessible in real life applications than the history-dependent ones (which require access to very detailed information).

long-run.⁵⁵ The logic of the overturning principle does not apply in the ABEE because agents do not make proper inferences from the timing of the decisions. In Guarino and Jehiel (2013), we have also considered the setting of Lee (1993) in which actions vary in the continuum, and obtain there with the same payoff-relevant analogy partition that in the long-run there would be no inefficiencies, despite the excessive weight attached to early signals.⁵⁶

Auctions

I have developed three insights using ABEE in auction environments. First, in private value environments with independent distributions of valuations across bidders, I have considered asymmetric First-Price Auctions with bidder-anonymous feedback partition (see Jehiel (2011)). I note that such an analogy partition can be motivated by the common practice of auction sites to disclose past bids but not bidders' characteristics. In the case of two bidders, I have shown that ABEE results in an efficient allocation and that the revenue is larger than that generated by the Second-Price Auction.⁵⁷ Such an insight reveals that by making the feedback about past bids anonymous, it is possible to do better than the Second-Price Auction from a revenue perspective without sacrificing on the welfare efficiency. Such an insight sharply contrasts with a traditional insight obtained with full rationality, suggesting a necessary trade-off between efficiency and revenue (as follows from the celebrated allocation equivalence result holding when the distributions of valuations are independent across bidders, as assumed here).

Second, Jehiel and Lamy (2015) have considered competing auction environments in private value environments. When buyers form their expectation about others' particip-

⁵⁵To get a sense of the result, observe that if there were no inefficiencies in the long-run, past observations in the sequence would be almost perfectly informative of the state, implying that agents except in the first period would not make use of their information. But, such behaviors would lead to long-run inefficiencies, as in the classic social learning model, leading to a contradiction.

⁵⁶Our insight in Lee's model is to be contrasted with that obtained by Eyster and Rabin (2010) who consider a model that mixes elements of cursedness (see below) with elements of subjective priors (in their model, subjects subjectively believe that their predecessors are more cursed than they really are). See also Bohren (2016) and Bohren and Hauser (2019) for further modelling of this.

⁵⁷Efficiency immediately follows from the bidder-anonymous feedback partition, as such an analogy partition implies that a bid is a sole function of the valuation and not of the identity of the bidder. That revenue is increased in the case of two bidders follows from the observation that it makes every bidder feels he is competing with a bidder of average strength, which in turn yields a strong bidder to bid more aggressively than if he knew he were competing with a weak bidder.

ation in the various auctions in a coarse way, i.e., not seeing the link between the reserve price and the participation decisions, competition between (rational) sellers leads them to set zero reserve price (or absolute auctions) even if their value for keeping the good is strictly positive. This follows from a logic similar to that in Bertrand competition, and differs sharply from the traditional insight in competing auction environments which would predict that the reserve price should be set at the value of the seller (see, for example, McAfee (1993)). Jehiel and Lamy (2015) also provide a rationale for the use of secret reserve price auctions when some share of buyers aggregate the reserve price choices across situations where they are kept secret and situations where they are publicly disclosed (thereby breaking the logic of the unravelling argument that prevails in the standard rationality case and that makes secret reserve price auctions hard to rationalize). To the extent that absolute auctions and secret reserve prices are often used, the ABEE methodology provides a rationale for such auctions formats that was missing with the standard rationality approach.

Third, Jehiel and Mierendorff (2021) consider private value auctions with correlation in which at the time of the auction bidders receive noisy signals about their ex-post value. It seems then plausible that when looking at past auctions, such signals would not be accessible and only the bids and the ex-post values would. This suggests that at least some bidders would use the ex-post payoff-relevant analogy partition (relating the distribution of others' bids to their own ex post values). In the case of correlation, such coarse analogy partitions lead bidders to reason as in interdependent value models, and when there are both coarse and rational bidders, inefficiencies then arise in the Second-Price auctions as well as in every auction-like mechanism.

Strategic communication

Several insights mixing ABEE with strategic communication setups have been developed. First, in cheap talk environments à la Crawford and Sobel (1982), Jehiel and Koessler (2008) have considered a setup in which to interpret the communication strategy of the Sender, the Receiver partitions the state space into two intervals. In some cases, the resulting ABEE induces more communication than in the standard rationality case.

Second, Ettinger and Jehiel (2021) have considered a repeated expert/agent relationship à la Sobel (1985) in which one of the periods has more weight and the expert may

either be rational or else be committed to always telling the truth (about a state that the expert is informed about and that changes in every period). When the rational expert and the agent have opposite preferences, and the agent reasons coarsely based on the aggregate behavior of the two types of experts across all periods, the ABEE is such that the rational expert should tell the truth up to the key period, and lie at the key period (a behavior we refer to as a deceptive tactic) so as to take advantage of the inflated belief of the coarse agent that the expert must be a truth-teller. This prediction differs from that in the Sequential Equilibrium (analyzed by Sobel) in which the Sender should be mixing and the deceptive tactic should not be more rewarding than any other chosen tactic (due to the indifference implied by the equilibrium mixing).⁵⁸

Third, I have considered cheap talk environments in which a lying Sender would not remember his past lies, and I have modeled the expectation of a forgetful liar by requiring that such a Sender would conjecture he previously sent a message according to the aggregate distribution of lies (that is, aggregating all states where there is a lie in a single analogy class to form his expectation, see Jehiel (2021)).⁵⁹ While in one-shot communication protocols, such an asymmetric memory between liars and truth-tellers would make no difference relative to the standard analysis, in situations in which the Sender is asked several times to report what he knows, it can be shown that the only possible ABEE require that there is almost perfect elicitation of what the Sender knows in the limit when the grid of possible states gets finer and finer. I note that a similar insight would not hold if liars were assumed to know how the distribution of lies depends on the state (i.e., removing the aggregation idea), and I use the obtained insights to suggest a novel rationale to practices commonly used in criminal investigations that consist in asking suspects similar questions several times.⁶⁰

⁵⁸The main object of Ettinger and Jehiel (2020) concerns an experiment. In the experimental data, while we do not see all senders employing the deceptive tactic, we observe that it is chosen with a probability much above that predicted by the Sequential Equilibrium and that it is more rewarding on average than the other chosen strategies. We interpret the heterogeneity of the data by assuming that a share of the subjects is coarse while the rest is rational.

⁵⁹To model such imperfect memory with ABEE reasoning, I model the Sender at two different times, as different players sharing the same objective. This is similar in spirit to the multi-self modeling in Piccione and Rubinstein (1997).

⁶⁰There are other uses of ABEE in strategic communication applications. These include applications to signalling (Mullainathan et al. (2008), Hagenbach and Koessler (2017) or Bilancini and Boncinelli (2018)) as well as the modeling of coarse language on the Sender's side in Hagenbach and Koessler (2019)

Mechanism design

In Jehiel (2011), I have considered mechanism design environments in which in addition to the usual auction rule instruments, the designer can decide to disclose the feedback about past bids in a more or less disaggregated way. For example, the designer can employ different auction formats in various proportions and disclose past bids as aggregated over different formats, or she can disclose feedback in an anonymous way, as discussed above in the context of asymmetric first-price auctions. In other words, through her choice of feedback policy, the designer is viewed as choosing the analogy partitions used by the bidders. A key insight is that the designer can always achieve a revenue that is strictly larger than in Myerson (1981)'s optimal auction by an appropriate choice of auction formats and an appropriate choice of feedback policy.⁶¹ This suggests a new role for the choice of feedback about past auctions that is typically not discussed in traditional mechanism design.

From a different perspective, in Jehiel (2019), I have suggested revisiting mechanism design questions assuming that at least some share of agents would rely on the payoff-relevant analogy partition, so as to reflect the constraint that it is generally hard to have access to the beliefs of agents at the time they submitted their report in the mechanism. The idea that beliefs may be hard to access is also at the heart of the motivation for the robust mechanism design literature (Bergemann and Morris (2005)). This has led a number of authors to consider stronger notions of implementation, such as ex-post implementation, which would work well irrespective of the beliefs of agents. Unfortunately, as shown in Jehiel et al. (2006), this is a very strong notion of implementation, so strong that in rich enough (i.e., multidimensional) environments it is hopeless to rely on it. An alternative could be to rely on an ABEE setting in which some share of agents would rely on the payoff-relevant analogy partition. Jehiel and Mierendorff (2021) can be viewed as an illustration of this where the notion of payoff-relevance is understood there at the

or on the Receiver's side (Eliaz et al. (2019)) who develop a model in which the Sender can choose the analogy partition of the Receiver as in Jehiel (2011)).

⁶¹To make the study of the optimal design interesting, one needs to place some constraints on the mechanisms such as those resulting from the ability of the buyers to reject the deal ex post. This is so because otherwise, the seller would be able to make infinite amount of money with risk neutral buyers, by (mis)leading buyers into having incorrect views on the distribution of some variables (as a result of their analogical reasoning) and then using betting devices related to these variables.

ex-post stage.

Investment and selection neglect

As already noted, ABEE can be used to model selection neglect. I have considered such a use of ABEE in the context of investment decisions (see Jehiel (2018)). In that application, every investor observes a noisy signal about his project and he has to decide whether or not to implement it. Investors know the cost but not the potential benefit associated to the project. To form his expectation about the mapping between the signal and the distribution of benefits, the investor considers those projects that have been implemented by others and for which he obtains the same signal. He extrapolates as if the obtained sample were representative. Accordingly, in the risk neutral case, he invests when the average benefit in the sample exceeds the cost. When the signals obtained by various investors about the same project are always the same, such a heuristic does not result in mistakes. But, when the signals may vary across investors (say because each economic agent adds an idiosyncratic error of his own when generating information on projects), such a selection neglect results in a bias. When the signal and benefit satisfy a Monotone Likelihood Ratio Property, the mistake results in an overoptimism bias with excessive investment in equilibrium. Moreover, the bias obtained in ABEE for those selection neglect investors is largest when the informativeness of signals is intermediate or when there are more rational investors around. Compared to models of overconfidence with subjective priors (see discussion of those below), the ABEE approach allows for an endogenous analysis of the effects of the underlying signal technology and/or of the share of sophisticated agents on the magnitude of the overoptimism bias. Such predictions have been confirmed in an experimental setting in Barron et al. (2020).

There have been other models related to selection neglect (see in particular, Esponda (2008) and Esponda and Vespa (2014)). I will come back to the concept introduced in Esponda (2008) later.

Some other applications

There have been a number of other applications of ABEE in different subfields. First, Jehiel and Koessler (2008) have revisited the email game studied by Rubinstein (1989) with the payoff-relevant analogy partition to establish that more equilibrium outcomes could then be supported. Second, there have been a number of finance applications in

which traders hold different theories as parameterized by how they categorize Markovian states into different analogy partitions (Eyster and Piccione (2013) in a setup à la Harrison and Kreps (1978) or Mailath and Samuelson (2020) for a study of whether traders with different theories or analogy partitions can disagree in the long-run, see also Piccione and Rubinstein (2003)). Third, there have been some applications to contractual relationships (see the working paper version of Jehiel (2015) in which coarse analogy partitions are considered in moral hazard in team interactions and Schumacher and Thysen (2020) for a study of optimal contracts with boundedly rational agents with a model that employs the Bayesian network machinery of Spiegel (2016)). Fourth, there have been some public economics applications in which agents may not see properly the effect of the change of some policy instruments (such as taxes) on other variables of interest (such as public goods) (see Grimee (2019) or Dal Bo et al. (2018) for an experimental study). Finally, Antler and Bachi (2020) consider a matching application of ABEE in a dynamic setting in which agents eventually remain unmatched due to their coarse analogy partitions used to assess the probability of acceptance of the other side of the market.

4 Related approaches

In this Section, I discuss how ABEE relates to other equilibrium concepts. I start the discussion with two classic concepts, the Bayes Nash equilibrium with subjective prior and the self-confirming equilibrium. I then discuss three approaches closely related to ABEE before opening the discussion to other equilibrium approaches to bounded rationality in games.

4.1 Pre-existing approaches

Two important notions of equilibrium that depart from the standard Nash equilibrium with objective prior have been proposed early in the literature. These are the Nash equilibrium with subjective priors proposed by Harsanyi (1967) and the self-confirming equilibrium that has appeared in various forms (see in particular Battigalli (1987), Fuden-

berg and Levine (1993) and Dekel et al. (2004)).⁶² I will now describe these equilibrium notions in the context of two-player static Bayesian games, and discuss how they relate to ABEE.

Nash equilibrium with subjective priors

Consider the following strategic situation. Player i 's utility is $u_i(a_i, a_{-i}; \theta_i, \theta_{-i})$ where $-i$ refers to the player other than i , $a_j \in A_j$ refers to j 's action, and $\theta_j \in \Theta_j$ refers to j 's type for $j = i, -i$. A Bayesian game also specifies a distribution $p(\theta_i, \theta_{-i})$ over types, which I will interpret as an objective distribution (as chosen by Nature). When this distribution is commonly known, a Bayes Nash equilibrium σ requires that if player i with type θ_i chooses a_i^* with positive probability, or in mathematical symbols if $\sigma_{\theta_i}[a_i^*] > 0$, then

$$a_i^* \in \arg \max_{a_i} \sum_{a_{-i}, \theta_{-i}} p(\theta_{-i} | \theta_i) \sigma_{\theta_{-i}}[a_{-i}] u(a_i, a_{-i}; \theta_i, \theta_{-i}).$$

That is, σ_{θ_i} should be a best-response to the correct view about how type θ_{-i} is distributed given θ_i (as reflected by the conditional distribution $p(\cdot | \theta_i)$) and the correct view about the probability with which player $-i$ with type θ_{-i} chooses the various actions a_{-i} (as given by $\sigma_{\theta_{-i}}[a_{-i}]$).

Harsanyi (1967) proposed a solution concept in which player i would be allowed to have his own subjective view about how θ_{-i} is distributed as a function of θ_i . Let $\hat{p}^i(\cdot | \theta_i)$ denote this subjective distribution. The corresponding solution concept called Bayes Nash equilibrium with subjective prior requires that player i best-responds to such a subjective perception about how types are distributed together with the correct perception about how the opponent behaves as a function of his type. Formally,

Definition 5 *A Bayes Nash equilibrium with subjective priors $\hat{p} = (\hat{p}^i)_{i=1,2}$ is a strategy profile σ such that, for every θ_i , if $\sigma_{\theta_i}[a_i^*] > 0$ then*

$$a_i^* \in \arg \max_{a_i} \sum_{a_{-i}, \theta_{-i}} \hat{p}^i(\theta_{-i} | \theta_i) \sigma_{\theta_{-i}}[a_{-i}] u(a_i, a_{-i}; \theta_i, \theta_{-i}).$$

While appealing in some respects, this solution concept is problematic in others. First,

⁶²The works of Kalai and Lehrer (1993) and Kurz (1994) can also be related to the self-confirming equilibrium.

it suggests an asymmetric treatment in the expectation formation about the strategy of nature (in charge of choosing the distribution of types) and the strategy of the opponent (in charge of choosing the distributions of actions as a function of the type). While arbitrary subjective beliefs are allowed for the former, correct beliefs are required for the latter. Presumably, the idea behind this concept is to allow for subjective beliefs about the strategy of nature in the tradition of Savage (1954) while sticking to the discipline device imposed by the correct expectation assumption in Nash equilibrium. However, it is somewhat unclear how to interpret such a concept from a learning perspective. If there are enough data to assess properly how the opponent behaves as a function of his type, it is then likely that there are also enough data to assess properly the strategy of nature.⁶³ But then, one should consider Bayes Nash equilibria with common objective priors and not with subjective priors. A related concern is raised in Dekel et al. (2004).

This is to be contrasted with the ABEE concept. In ABEE, players have a subjective belief about the play of other players (possibly including nature, see some of the above applications in which Nature is treated as a player), and these subjective beliefs are structured through the consistency requirement by the equilibrium play together with how players bundle contingencies into analogy classes. This in turn explains why ABEE can be interpreted as a limiting outcome of a learning process with coarse feedback (see subsection 2.3 above for a sketch of learning model supporting this view), while such a rationale is not readily available for the Bayes Nash equilibrium with subjective priors.

Another concern about the Bayes Nash equilibrium with subjective prior is that it is silent on how players choose their subjective priors. Without guidance on this, the approach is not very predictive, as many different behaviors can be rationalized by varying the subjective priors.⁶⁴ Of course, one may use some of the insights derived experimentally

⁶³Think of a situation in which (a, θ) would be observed after the play of a game so as to allow to assess correctly how a_{-i} is distributed for each θ_{-i} . Such dataset would also allow to assess correctly how θ_{-i} is distributed for each θ_i , making the subjective view difficult to motivate. And even if player i cannot recover how (θ_i, θ_{-i}) is distributed, I would say the choice of subjective prior \hat{p}_i is somehow arbitrary, as it is not guided by any objective information available to i .

⁶⁴A similar concern can be raised about ABEE, as one considers the span of behaviors that can be rationalized when varying the choices of analogy partitions. I see several possible responses to such a concern in the ABEE context. First, when the analogy partitions are viewed as resulting from the objective feedback subjects are exposed to, these can be considered as part of the description of the interaction in the same way as information partitions are (and there is no reason to worry about which predictions would arise with other analogy partitions). Second, assuming it is not known which aggregate

in behavioral studies to suggest which kinds of subjective priors would be more relevant. For example, biases such as overconfidence may be modelled by requiring that players think their private information is more precise (about the state of the economy) than it really is. A number of influential models in behavioral finance rely on such a modeling (see, for example, Scheinkman and Xiong (2003)). Observe that this view of overconfidence should be contrasted with the overoptimism derived in the investment decision problem considered in Jehiel (2018) in which the overoptimism is endogenously shaped by the objective statistical structure relating the variables of interest and the presumption that decision makers extrapolate from the sole observations of past implemented projects (as captured by appropriately defined analogy partitions, see the above discussion of selection neglect).

Self-confirming equilibrium with general feedback structure

I describe a formulation of self-confirming equilibrium somewhat similar to that in Dekel et al. (2004) which I adapt to capture that I have in mind that a given subject in the role of a player would play just once. The strategic situation is still that of a static Bayesian game with two players i and $-i$ with an objective distribution of types as given by $p(\cdot, \cdot)$. Such a game is played many times by pairs of subjects $(i, -i)$ randomly matched from the population of players i and $-i$. After a game is played, some signals y_i and y_{-i} are made available to subjects in the population of players i and $-i$, respectively. Formally, if the played action profile is $a = (a_i, a_{-i})$ and the profile of types is $\theta = (\theta_i, \theta_{-i})$, subjects in the population of player i receive a signal y_i assumed to take finitely many values in Y_i where y_i is drawn from a distribution $\tilde{y}_i(a; \theta)$ that may depend on $(a; \theta)$. Such a signal y_i may include information on i 's obtained payoff (as almost always assumed in the self-confirming literature where the same player i is viewed as playing many times) or any other element that would be informative of some characteristics of $(a_i, a_{-i}, \theta_i, \theta_{-i})$.

A self-confirming equilibrium aims at capturing situations in which player i could hold subjective theories both about how nature chooses types and how players behave as

feedback players are exposed to, studying which insights are robust to various analogy partitions seems less demanding than studying which insights are robust to various subjective priors, at least when the number of contingencies is not too large. While not much has been done in relation to robustness in the ABEE setting, I certainly regard the robustness question as a valid research question worthy extra work. See also the final part of the paper for further elaborations on the choice of analogy partitions.

a function of their type where the subjective theories of player i are required to give rise to a distribution of y_i that matches the observed frequencies. Formally,

Definition 6 *A strategy profile σ is a self-confirming equilibrium (given $\tilde{y}_i, \tilde{y}_{-i}$) if, for every player i , there exist $\hat{p}^i, \hat{\sigma}^i = (\hat{\sigma}_{\theta_{-i}}^i, \hat{\sigma}_{\theta_i}^i)$ such that*

- 1) *If $\sigma_{\theta_i}[a_i^*] > 0$ then $a_i^* \in \arg \max_{a_i} \sum_{a_{-i}, \theta_{-i}} \hat{p}^i(\theta_{-i} | \theta_i) \hat{\sigma}_{\theta_{-i}}^i[a_{-i}] u(a_i, a_{-i}; \theta_i, \theta_{-i})$ and*
- 2) *For every $\bar{y}_i \in Y_i$, $\Pr_{p, \sigma}(y_i = \bar{y}_i) = \Pr_{\hat{p}^i, \hat{\sigma}^i}(y_i = \bar{y}_i)$.*

A self-confirming equilibrium can be viewed as a steady state of a dynamic process in which subjects in the role of player i would observe signals y_i from previous interactions, and form their subjective theories based on these observations. While the first condition in Definition 6 ensures that player i chooses a best-response to his subjective theories, the second condition ensures that the subjective theories of player i induce the same distribution of signals y_i as the distribution that is truly observed, which is generated by the true p and the true σ (which player i is assumed not to know). Another way to think of the second condition is that if it were not satisfied, the data observed by subjects in the population of player i would allow them to conclude that their subjective theories cannot be true. It is called self-confirming equilibrium in recognition of the property that the observed data do not refute (or less accurately confirm) the subjective theories.

I suggest several comments in relation to Definition 6. First, a Bayes Nash equilibrium (with objective prior p) is always a self-confirming equilibrium, no matter what the signal specification $y = (y_i, y_{-i})$ is.^{65,66} Of course, for many signal specifications y_i , it is unclear how player i would be able to have the correct theory (think of the case in which y_i would be totally uninformative). But, the definition of self-confirming equilibrium makes no restriction on the subjective theories, as long as they are not refuted by the observations. Since the correct theories are obviously not refuted, it follows that a Bayes Nash equilibrium is always a self-confirming equilibrium (even when there is no clear way/procedure

⁶⁵This follows because if $\hat{p}^i = p$ and $\hat{\sigma}^i = \sigma$, the second condition is trivially satisfied. And replacing \hat{p}^i by p on the one hand and $\hat{\sigma}^i$ by σ on the other in the first condition directly implies the condition for a Bayes Nash equilibrium.

⁶⁶This has led game theorists to characterize the conditions on y that would ensure that the set of self-confirming equilibria coincides with the set of Bayes Nash equilibria. For example, if $\tilde{y}_i(a; \theta)$ completely identifies $(a; \theta)$, all self-confirming equilibria are Bayes Nash equilibria. This is also so in a private value settings, if $\tilde{y}_i(a; \theta)$ identifies (a_{-i}, θ_i) so that i can properly learn the distribution of a_{-i} conditional on his type θ_i .

for the players to guess the correct theories from the feedback they receive). In contexts in which the full description of the game is common knowledge to the players, it is possible to refine the subjective theories to only allow those that can be justified by some rational behavior. This leads to the rationalizable self-confirming equilibria as defined in Rubinstein and Wolinsky (1994).⁶⁷ For reasons similar to those just expressed, note that all Bayes Nash equilibria (with objective prior p) are always rationalizable self-confirming equilibria whatever the signal specification y .

A second comment is that I have allowed in Definition 6 player i to form subjective theories about the play of other subjects in the population assigned to the role of player i . While this subjective theory plays no role in the first (best-response) condition, it may in the second, as the behavior of other subjects in the role of player i may affect the distribution of observed signals. This is somewhat unusual in the literature on self-confirming equilibrium in which it is generally assumed that the same subject would be in the role of player i , but I would argue it is natural once one has in mind situations in which different subjects would play the role of player i .⁶⁸

A third comment is that Definition 6 assumes that all subjects in the role of player i have the same subjective theories. This is referred to as the unitary case in the self-confirming equilibrium literature (see Fudenberg and Levine, 1998) and can be generalized to allow different subjects in the role of player i to entertain heterogeneous subjective theories, all assumed to satisfy the second (consistency) condition.

It should be mentioned here that the analogy-based expectation equilibrium in Bayesian games as considered in Jehiel and Koessler (2008) can be related to the self-confirming equilibrium for a particular class of signal structure y . Indeed, considering a Bayesian game with analogy partitions defined over the state space of Θ , one can view an ABEE in such a setting as a self-confirming equilibrium in which after a play $(a; \theta)$, a subject in the population of players i receives the signal $y_i(a; \theta) = (a_{-i}, \alpha_i(\theta))$ where $\alpha_i(\theta)$ denotes the analogy class of player i to which θ belongs. This is so because the signals y_i allow subjects in the role of player i to recover the aggregate distribution of player

⁶⁷See Dekel et al. (1999) for a related concept in extensive form games when players observe the equilibrium path.

⁶⁸A similar view (with players forming conjectures about players assigned to the same role) appears in Fudenberg and Kamada (2018).

$-i$'s action in each analogy class α_i (corresponding to $\beta(\alpha_i)$ in the ABEE setting), and obviously the subjective theory that views the distribution of behavior of $-i$ at θ as coinciding with $\beta(\alpha_i(\theta))$ automatically satisfies the second requirement in the definition of a self-confirming equilibrium.

From this perspective, a key distinctive feature of the ABEE is that the subjective theory used in ABEE is not any theory compatible with the feedback, but it is a special one derived from the feedback in a trivial way (very similar to the identity mapping), thereby making the mapping from the signal specification y_i to the subjective theory more transparent than in the self-confirming equilibrium. An implication of this is that the ABEE corresponds to a refinement of self-confirming equilibrium, and it is no longer the case that a Nash equilibrium is necessarily an ABEE (a conclusion that cannot be obtained with the self-confirming equilibrium). As already highlighted, the selection principle in ABEE can be related to the central theme of simplicity in bounded rationality: the subjective theory considered by player i is the simplest one that is compatible with the observed feedback y_i (see discussion above). While the feedback considered for ABEE allows for a natural simplicity-based selection of the subjective theories, an open question is whether a similar approach can be proposed for other signal specifications in an insightful way.

More recently, Battigalli et al. (2015) have proposed extending the self-confirming equilibrium to allow players to exhibit ambiguity aversion over the various possible subjective theories compatible with the feedback y (see also Lehrer (2012)). While their most general approach always includes the traditional self-confirming equilibria as possible outcomes, one idea discussed in their paper is the maxmin self-confirming equilibrium in which the worst case criterion is applied to the entire set of all subjective theories compatible with the feedback. The maxmin self-confirming equilibrium so defined is a refinement of the self-confirming equilibrium. Such a refinement can be viewed as an alternative to ABEE. I see however several issues with the general use of this refinement in applications. First, given that in the proposed maxmin approach, all compatible theories are considered, the maxmin criterion should be regarded as modelling an extreme form of ambiguity aversion, and there is no reason in general to believe that all agents would have such an extreme attitude toward ambiguity (in decision contexts, this has led Klibanoff et al. (2005) to consider smoother versions of ambiguity aversion, which is the

main modeling approach in Battigalli et al (2015)). Second, even accepting this extreme form of ambiguity aversion, the mere determination of the set of all subjective theories that would be compatible with the observed feedback may be quite challenging (not to mention the complexity of determining the maxmin computed on the basis of this set). Given this, it is unclear to me whether such a selection can capture the behavior of real (boundedly rational) economic agents.⁶⁹

4.2 Closely related approaches

In this part, I discuss the cursed equilibrium introduced by Eyster and Rabin (2005) contemporaneously to the ABEE concept, as well as the Bayesian network equilibrium and the Berk-Nash equilibrium introduced later by Spiegel (2016) and Esponda and Pouzo (2016), respectively. These concepts are closely related to the ABEE, as I will explain.

Cursed equilibrium

Consider Bayesian games as defined in the previous subsection. Given a strategy profile σ and a type θ_i of player i , define $\bar{\sigma}_{-i}^{\theta_i}$ as the type-independent strategy of player $-i$ satisfying

$$\bar{\sigma}_{-i}^{\theta_i}[a_{-i}] = \sum_{\theta_{-i}} p(\theta_{-i} | \theta_i) \sigma_{\theta_{-i}}[a_{-i}].$$

It is the strategy of player $-i$ aggregated over the realizations of θ_{-i} conditional on the realization θ_i of i 's type.

The partially cursed equilibrium of Eyster and Rabin assumes that player i with type θ_i best-responds to the belief that player $-i$ plays a strategy that is a convex combination between $\bar{\sigma}_{-i}^{\theta_i}$ and the true strategy σ_{-i} . Formally, for each χ , σ_{-i} , θ_i , define $\sigma_{\theta_{-i}}^{\chi, \theta_i}$ such that

$$\sigma_{\theta_{-i}}^{\chi, \theta_i}[a_{-i}] = \chi \bar{\sigma}_{-i}^{\theta_i}[a_{-i}] + (1 - \chi) \sigma_{\theta_{-i}}[a_{-i}].$$

⁶⁹Of course, only empirical/experimental studies can address this question in a fully satisfactory manner. As a possible partial exploration of this, note that in the context of the centipede game, providing feedback on the aggregate pass rate as in Danz et al. (2016), the maxmin self-confirming equilibrium would coincide with the Nash equilibrium. Experimental data suggest otherwise (they are more in line with the ABEE, see Danz et al (2016) and Jehiel (2005)).

Definition 7 A χ -cursed equilibrium is a strategy profile σ such that for every player i and every type θ_i , if $\sigma_i[a_i^*] > 0$ then

$$a_i^* \in \arg \max_{a_i} \sum_{a_{-i}, \theta_{-i}} p(\theta_{-i} | \theta_i) \sigma_{\theta_{-i}}^{\chi, \theta_i}[a_{-i}] u(a_i, a_{-i}; \theta_i, \theta_{-i}).$$

Several comments are in order. First, when $\chi = 0$, a cursed equilibrium coincides with a Bayes Nash equilibrium. Second, when $\chi = 1$, the cursed equilibrium referred to as a fully cursed equilibrium by Eyster and Rabin coincides with an ABEE in which the analogy partition of player i coincides with his private information. This can be seen as the aggregate distribution of a_{-i} over all states $\theta = (\theta_i, \theta_{-i})$ with the same player i 's type θ_i precisely coincides with $\bar{\sigma}_{-i}^{\theta_i}$, and thus the best-response requirement in the fully cursed equilibrium is exactly the same as that for an ABEE with private information analogy partitions in this class of Bayesian games.

Considering interior values of χ allows Eyster and Rabin to go smoothly from the case of full sophistication as arising in a Bayes Nash Equilibrium (with $\chi = 0$) to a case of full naivete (with $\chi = 1$) viewed here as the failure to appreciate the link of others' strategy to their private information (which matters for interdependent value contexts for which this concept was introduced). Such a modelling of partial sophistication differs from that in ABEE in which the degree of sophistication is parameterized by the coarseness of the analogy partition. An advantage of Eyster and Rabin's approach compared to ABEE is that it is more parsimonious (even if by its focus on Bayesian games with interdependent values, it is also more limited in scope).⁷⁰ But, this advantage comes at a cost, as Eyster and Rabin do not provide a rationale for how player i with type θ_i would be led to have a belief about player $-i$'s strategy as given by $\sigma_{\theta_{-i}}^{\chi, \theta_i}$ when χ lies in the interior of $[0, 1]$. Clearly, if player i need to know separately $\sigma_{\theta_{-i}}$ and $\bar{\sigma}_{-i}^{\theta_i}$ to determine $\sigma_{\theta_{-i}}^{\chi, \theta_i}$, this is not satisfactory, as the separate knowledge of $\sigma_{\theta_{-i}}$ should allow player i to determine the best-

⁷⁰I would like to add that in my view parsimony alone is not a good criterion to select a theory or model (see also Spiegel (2019) who expresses a similar view). An analogy with the classic treatment of information in games may be instructive. In the spirit of the partially cursed equilibrium, one might have modeled partial information by requiring that a player plays a convex combination between what is optimal to do for him when he is fully informed of the state and what is optimal for him when he is totally uninformed. This is not how partial information is modelled though, and a partitional approach is generally adopted instead. I would say this is so because the partitional approach but not the convex combination approach can receive a natural interpretation in this case.

response to the actual behavior of player $-i$. But, then we are left to find an alternative argument as to how player i would have a belief given by $\sigma_{\theta_{-i}}^{\chi, \theta_i}$ when there is no knowledge of $\sigma_{\theta_{-i}}$. One possible idea in the vein of the self-confirming equilibrium would be to assume that player i observes from past interactions the aggregate behavior of player $-i$ for each possible θ_i as in the motivation for the private information analogy partition. The χ -cursed equilibrium is a self-confirming equilibrium with this signal structure. However, while $\sigma_{\theta_{-i}}^{\chi, \theta_i}$ is a possible subjective theory that would give rise to what player i observes from past interactions, I would argue it is unclear how player i would get at such a subjective theory (given that i has no knowledge of $\sigma_{\theta_{-i}}$, and only $\bar{\sigma}_{-i}^{\theta_i}$ is accessible from the data coming from past interactions). In my view, a satisfactory rationale for the partially cursed equilibrium is missing.⁷¹ Some readers may wonder whether it is not possible to interpret the χ -cursed equilibrium as an ABEE in which a share χ of players would use the private information analogy partition and a share $1 - \chi$ of players would use the finest analogy partition.⁷² However, such an interpretation is incorrect, as in a χ -cursed equilibrium, player i with type θ_i always holds the same belief which is $\sigma_{\theta_{-i}}^{\chi, \theta_i}$ whereas in the ABEE proposed alternative, a fraction χ would hold the belief $\bar{\sigma}_{-i}^{\theta_i}$ and a fraction $1 - \chi$ would hold the belief $\sigma_{\theta_{-i}}$ (resulting in general in different best-responses).

Bayesian network equilibrium

The Bayesian network equilibrium introduced by Spiegler (2016) can be described as follows. Inspired by the work on causal networks of Pearl (2009), Spiegler proposes endowing decision makers (or more generally players) with subjective causality binary

⁷¹Miettinen (2009) suggests an interesting way to view a partially cursed equilibrium as an ABEE of a different Bayesian game in which the state space would be enlarged. However, the construction does not relate to a natural feedback structure, making such a construction more of theoretical than practical interest.

Siegler (2020) suggests a different interpretation of the partially cursed equilibrium in which players would be exposed to two sets of statistics, the one corresponding to the private information analogy partition and the one corresponding to the fine analogy partition, and players would mix the two theories using a convex combination of the corresponding expectations. I will investigate at the end how a player exposed to different coarse statistics corresponding to different analogy partitions could behave. While the mixture corresponding to the partially cursed equilibrium is a possibility, I suggest several others as well. Moreover, it is unclear why the two coarse statistics needed for the partially cursed equilibrium should naturally be assumed to be available to players in applications, making this rationale a bit weak for applied purposes.

⁷²If this were so, we would have a rationale for the χ -cursed equilibrium similar to that for ABEE with heterogeneous analogy partitions

relations that specify which variable is caused by which other variables. Formally, such a causality network is represented by a Direct Acyclic Graph (DAG) relating the various variables of interest x_0, x_1, \dots, x_K . When $i \rightarrow j$ in the DAG, it means that the decision maker thinks that (the distribution of) x_j is influenced by the realization of x_i , thereby formalizing the intuitive idea that x_i is a cause of x_j .⁷³ From a game-theoretic viewpoint, one of the most interesting aspects in Spiegler (2016)'s approach is the equilibrium nature of the concept, which is due to the endogenous character of one variable, the action a of the decision maker.

To present formally the concept, let $x = (x_0, x_1, \dots, x_K)$ be the variables of interest where $x_0 = a$ denotes the action to be chosen by the decision maker. Given the DAG considered by the decision maker, one can define for each i , the set of j such that there is a link $j \rightarrow i$. Call this set $R(i) \subset K = \{0, 1, \dots, K\}$. The variables x_k , $k \in R(i)$ and only those are viewed by the decision maker as causing x_i .

To simplify the presentation, I will be assuming that the decision maker receives no private information before making his choice of action. Importantly, I will also be assuming that 0 (corresponding to the decision maker's action) is the ancestor of the subjective DAG considered by the decision maker. I would argue this can be viewed as deriving from logical implications, as the decision maker's choice of action can hardly be thought of as being influenced by variables he does not observe. The true joint distribution of x will be given by $p(x)$ taking the form $p(x) = p(a)p(x_{-0} | a)$ where $x_{-0} = (x_1, \dots, x_K)$. The key object in a Bayesian network equilibrium is the subjective perception of how variables are distributed. It is given by the factorization formula:

$$p^{DAG}(x) = \prod_{i \in K} p(x_i | x_{R(i)}) \quad (4)$$

where $x_{R(i)}$ denotes the profile of x_k for $k \in R(i)$. The interpretation of this formula is that from the observed data generated by p , the decision maker can construct the conditional distributions $p(x_i | x_{R(i)})$, and given that the decision maker thinks there are no other dependences beyond those described by DAG, p^{DAG} is viewed as (subjectively) describing

⁷³Moreover, as indicated in the term DAG, there is no cycle in the graph, i. e., one cannot have $x_1 \rightarrow x_2 \rightarrow \dots \rightarrow x_n \rightarrow x_1$ for any subset of variables (x_1, \dots, x_n) .

the data generating process. The decision maker then chooses a so as to maximize

$$\sum_x p^{DAG}(x_{-0} | a)u(a, x_{-0}) \quad (5)$$

where u is the von Neuman-Morgenstern utility as usually defined. In order to ensure that $p^{DAG}(x_{-0} | a)$ is well defined from p^{DAG} for all a , Spiegler considers a trembling-hand formulation in which all actions should be played with probability no smaller than ε . Formally,

Definition 8 *A fully mixed strategy $p(a)$ is a personal ε -perturbed equilibrium if whenever $p(a) > \varepsilon$, a maximizes (5). A Bayesian network equilibrium is the limit of a personal ε -perturbed equilibrium as ε tends to 0.*

I will now explain how a Bayesian network equilibrium can be viewed as an ABEE of an appropriately defined extensive form game with appropriately defined analogy partition. I would like to suggest that the link to ABEE should come as no surprise given that the conditional probabilities $p(x_i | x_{R(i)})$ can precisely be viewed as the aggregate frequencies of x_i in the pool of data in which $x_{R(i)}$ is kept the same. That is, viewing nature in charge of choosing x_i as a player, say player i , the causality idea behind $R(i)$ means that the decision maker considers the strategy of i to be the same whenever $x_{R(i)}$ is kept constant and is independent of other variables. The ABEE is well suited to model such (partial) independences. It requires though defining a multi-stage game in which x_i is viewed as being chosen by player i , as just explained, where player i is viewed as moving after player j if $j \in R(i)$. It also requires that another player viewed as irrelevant from the viewpoint of the objective of the decision maker, call it Nature, chooses x right after the choice a of the decision maker according to $p(\cdot | a)$, as objectively defined.⁷⁴ Player i down the tree is then required to choose an x_i that matches the initial choice made by Nature.⁷⁵ While the game tree so constructed is subjective, I would argue it is naturally derived from the

⁷⁴Observe that it is the same decomposition of Nature into several players as the one considered above to model correlation neglect.

⁷⁵I need such a decomposition (between Nature choosing x_{-0} right after a and then player i choosing x_i later on) to allow for the true generating process not to share the factorisation structure imposed by the subjective DAG.

DAG to the extent that a causality link $i \rightarrow j$ directly suggests that the causing variable x_i must be decided before the consequence x_j .

More precisely, consider a multi-player extensive form game with one different player for each variable x_k , $k \in K$. Call the player in charge of x_k with $k \neq 0$, player k . The decision maker is player 0, and he chooses $x_0 = a$ at the root of the extensive form game. Then Nature also viewed as a player decides how $\omega = (x_k)_{k \in K}$ is distributed as a function of a according to the objective conditional distribution $p(\omega | a)$. Then the extensive form game specifies that each player k plays in sequence in an order that respects the order in the DAG. That is, all players in $R(i)$ must move before i .⁷⁶ Player k chooses x_k and his strategy is exogenously specified to be $x_k = x_k(\omega)$ where $x_k(\omega)$ is the k^{th} element in ω as chosen by Nature after the action a of the decision maker. The objective of the decision maker is $u(a; x)$ where a is his action, $x = (x_k)_{k \in K \setminus \{0\}}$ is the profile of actions chosen by players $k \in K \setminus \{0\}$, and u is the von Neuman-Morgenstern utility function defined above.

Regarding the analogy partition of the decision maker, consider the following. Decision nodes of different players belong to different analogy classes. For every player $i \neq 0$, assume that all his nodes corresponding to the same $x_{R(i)}$ define one analogy class.^{77,78} By varying $i \neq 0$ and $x_{R(i)}$, one obtains an analogy partition.

It is readily verified that the ABEE of the strategic setting just defined are the Bayesian network equilibria, as defined in Spiegel (2016). This is so because condition (4) is equivalent to the consistency condition in ABEE, and the ensuing maximization of (5) corresponds to the best-response in ABEE.

This construction allows me to view BNE as a special case of ABEE.⁷⁹ As such, the various applications of BNE can also be viewed as applications of ABEE.⁸⁰ Importantly,

⁷⁶Observe that one can always find such an order due to the absence of cycles in a DAG. There may sometimes be multiple extensive form games that are compatible with the order in the DAG (think of the case where the DAG is empty for example in which case any order is possible), but any of those would lead to the same set of equilibrium outcomes.

⁷⁷Observe that this is well defined because the extensive form game is such that players in $R(i)$ play before i .

⁷⁸For Nature, the player moving second, note that the partitioning of its nodes is irrelevant (given that $u(a; x)$ does not depend on ω).

⁷⁹In fact, within the extensive form game just defined, analogy partitions not corresponding to BNE can be considered, for example putting together nodes of different players i and j . This suggests that the ABEE formulation allows for forms of bounded rationality not covered by BNE.

⁸⁰I would like to stress that the motivation of analogy partitions through causality is new to Spiegel

the construction requires that the DAG considered by the decision maker is such that $R(0)$ is empty, as the decision maker is assumed here not to have any private information.⁸¹ If this were not the case or more generally if the decision maker were allowed to think his action is caused by variables he does not observe (as allowed but not considered in applications by Spiegler (2016)-(2020)), then it would not be possible to make this construction (because the causality network would require that i plays before the decision maker when $i \rightarrow 0$ in the proposed construction, and the maximization of (5) implicitly requires that the decision maker thinks his choice of a affects x_i through the term $p^{DAG}(x_{-0} | a)$, which would not be so in the corresponding ABEE). As argued above, I think that the requirement that $R(0)$ is empty when the decision maker has no private information is natural, and dropping it would in my view require a more severe form of bounded rationality (involving some kind of failure of internal logic of the decision maker).

It should be mentioned that when Spiegler (2016)-(2020) discusses ABEE in relation to BNE, he refers to the ABEE formulation in static Bayesian games as developed in Jehiel and Koessler (2008). In such contexts, when applied to decision problems, ABEE cannot capture the equilibrium nature of the strategy $p(a)$ of the decision maker in BNE. Yet, as just discussed, the equilibrium nature of $p(a)$ is captured by ABEE when applied to the extensive form game, as defined above.

Berk-Nash equilibrium

Esponda and Pouzo (2016) present a general solution concept in games aimed at capturing limiting outcomes of learning processes in which players would have misspecified theories to represent their environment. Their solution concept is parameterized by the game, what is observed by the various players at the end of each game as well as the set of subjective theories each player considers. Their solution concept combines ideas from statistics (in particular as developed by Berk (1966)) with equilibrium ideas (as in the Nash equilibrium), hence the term Berk-Nash equilibrium.

To present succinctly the key ingredients of the concept, consider a class of normal

(2016) and does not appear in earlier work on ABEE (maybe with the slight exception of the monitoring game in Ettinger and Jehiel (2005) in which some causality interpretation -moral hazard vs adverse selection - is mentioned). As illustrated by the many insightful applications of BNE, bringing a new motivation broadens the scope of applications.

⁸¹More generally, $R(0)$ should contain at most those i such that the decision maker observes x_i when there is private information.

form games with two players $i = 1, 2$ in which player i chooses $a_i \in A_i$, player i observes $y_i \in Y_i$ after the game is played, and the stage game payoff of player i is $u_i(a, y_i)$ where $a = (a_i, a_{-i})$ is the action profile chosen by the players. The true distribution of y_i given a is governed by some objective distribution $p(y_i | a)$, which is not known to player i . Instead, player i is endowed with a set Φ_i of subjective theories ϕ_i where $\phi_i \in \Phi_i$ specifies a distribution $\phi_i[a_i]$ over (a_{-i}, y_i) for each $a_i \in A_i$. I will be assuming that A_i, Y_i, Φ_i are all finite sets.⁸² I do not specify the prior of player i over Φ_i , as this is not needed to define the Berk-Nash equilibrium. Importantly, Φ_i may be misspecified in the sense of not including the ϕ_i such that for every a_i , $\phi_i[a_i]$ would be the true distribution over (a_{-i}, y_i) as generated by p and the strategy σ_{-i} of player $-i$.

The general principle behind a Berk-Nash equilibrium is as follows. Assume players keep playing according to σ . Together with p , this induces for each $a_i \in A_i$, a true distribution of y_i after action a_i of player i . Denote the corresponding probability of y_i by $p_{a_i, \sigma}[y_i]$. Viewing player i as trying to learn which $\phi_i \in \Phi_i$ to consider would lead him to only keep those ϕ_i which minimize:⁸³

$$KL_i(\phi_i, \sigma) = \sum_{a_i} \sigma_i[a_i] \sum_{y_i} p_{a_i, \sigma}[y_i] \ln \frac{p_{a_i, \sigma}[y_i]}{\phi_i[a_i][y_i]}$$

where $\phi_i[a_i][y_i]$ denotes the probability that signal y_i arises when i plays a_i according to the distribution $\phi_i[a_i]$.

KL_i is the expected (with respect to a_i) Kullback-Leibler divergence from ϕ_i to $p_{\cdot, \sigma}$. The work of Berk (1966) ensures that if the true process is governed by $p_{\cdot, \sigma}$, only subjective theories in $\arg \min_{\phi_i \in \Phi_i} KL_i(\phi_i, \sigma)$ can be in the support of the long-run belief μ_i of player i , as these are the theories which have the largest likelihood in Φ_i given the long-run observations of a_i and y_i .

The other requirement in a Berk-Nash equilibrium is the familiar best-response idea: σ_i should be a best-response to player i 's belief about (a_{-i}, y_i) after each possible a_i where

⁸²Esponda and Pouzo (2016) consider the case when Φ_i have infinitely many elements so as to capture the case of parametric estimations.

⁸³This is assuming that i knows σ_i which fits applications in which the same player i keeps playing. Otherwise, ϕ_i could also specify a theory about how other players in the role of player i behave (see the above comments surrounding the definition of self-confirming equilibrium).

such a belief can be inferred from the probability μ_i assigned by player i to the various ϕ_i in Φ_i . Putting together these two conditions yield:⁸⁴

Definition 9 *A strategy profile σ is a Berk-Nash equilibrium if, for every player i , there exists a belief μ_i over Φ_i such that*

- 1) $\mu_i[\phi_i^*] > 0$ implies that $\phi_i^* \in \arg \min_{\phi_i \in \Phi_i} KL_i(\phi_i, \sigma)$ and
- 2) $\sigma_i[a_i^*] > 0$ implies that $a_i^* \in \arg \max_{a_i} \sum_{\phi_i} \mu_i[\phi_i] E_{a_{-i}, y_i}(u(a_i, a_{-i}, y_i) \mid a_i, \phi_i)$.

Under some extra assumptions, the existence of Berk-Nash equilibrium is guaranteed (see Esponda and Pouzo (2016)). Recent research shows that not all Berk-Nash equilibria can emerge as the limiting outcome of learning processes (see Fudenberg et al (2020)), but I will abstract in my discussion from the resulting refinements.

How does the Berk-Nash equilibrium relate to ABEE? A first observation is that it is possible to interpret the ABEE as a limiting outcome of a learning environment in which each player i would (subjectively) believe that, for every analogy class α_i in the analogy partition of player i , the behavior is the same for all elements belonging to α_i . Indeed, with such misspecifications, the observations of the behaviors of j at τh for the various $(j, \tau h) \in \alpha_i$ would lead, in the long run, player i to adopt the view, using the notation introduced for ABEE, that player j behaves according to $\hat{\sigma}_{-i}[c_i; \beta]$ where β should be consistent with σ (as such a subjective view would be what minimizes the Kullback-Leibler divergence from the admissible subjective theories of player i to the true long-run data generating process).

Note that with this interpretation, the signal y_i observed by player i at the end of a game should include a record of the actions observed in the various analogy classes, but it may include more information as well given that, in the Berk-Nash equilibrium approach, the belief $\hat{\sigma}_{-i}[c_i; \beta]$ is the result of the misspecified prior rather than the coarseness of the observation y_i .⁸⁵ This is different in spirit from how I interpreted ABEE above when I viewed the belief $\hat{\sigma}_{-i}[c_i; \beta]$ as a consequence of the coarseness of y_i (together with some simplicity considerations), and not as a consequence of an initial misspecification of the

⁸⁴I give here the definition without trembling. As for many other approaches, it is natural to refine the concept by imposing some (vanishing) trembles that allow player i to have observations for all actions and not only those in the support of σ_i .

⁸⁵In particular, even if player i observes finely the behavior of the various players at every node and for every type, the same Berk-Nash equilibrium would arise.

prior (see the above learning interpretation of ABEE and the discussion surrounding the self-confirming equilibrium).

More generally, Esponda and Pouzo (2016) take the subjective theories Φ_i of the various players i as a primitive that is unrelated to the underlying objective characteristics of the game (as represented by the utility, information and feedback y associated to each player). To make good use of their approach, it would be nice in my view to make progress on how the subjective theories considered by the players may be shaped by the objective characteristics of the environment.⁸⁶ To some extent, the ABEE provides an illustration of how the feedback structure y may shape the subjective priors for a special class of y . An open question concerns the possibility to go beyond for other classes of feedback.

4.3 Other approaches

I now discuss various equilibrium approaches with some bounded rationality element. I also include a discussion of Case-Based Reasoning and Learning Classifier Systems that share some common motivations with ABEE, even if not taking the form of equilibrium approaches.

Valuation equilibrium

Jehiel and Samet (2007) define a solution concept for extensive form games of complete information, the Valuation equilibrium, that can be viewed as the limiting outcome of a learning process in the tradition of reinforcement learning models in which instead of reinforcing strategies, players reinforce clusters of moves referred to as similarity classes. The valuation equilibrium (VE) can be viewed as the reinforcement learning counterpart of ABEE.

To give a quick description of the concept, consider an extensive form game of complete information with players $i = 1, \dots, n$ (and possibly nature) where Z is the set of terminal nodes, N is the set of non-terminal nodes, A is the set of arcs (n, m) in the tree, N_i is the set of nodes where player i must move, $M_i(n) = \{m \mid (n, m) \in A\}$ is the set of moves of

⁸⁶This is particularly relevant if one has in mind that it is not possible to have access (in the field) to the set of subjective theories of the players, since allowing for any possible such sets would make the predictive power of the approach less strong. Such a concern may be related to a concern I expressed above in relation to Bayes Nash equilibria with subjective priors.

player i at node $n \in N_i$, and $u_i : Z \rightarrow \mathfrak{R}$ is the von Neuman-Morgenstern payoff function defined over terminal nodes.

The key novel parameter of the strategic setting in VE is, for each i , the similarity partition Λ_i of player i , which is a partition of $M_i = \bigcup_{n \in N_i} M_i(n)$, the set of all moves that player i may choose in the game. For any $m \in M_i$, I let $\lambda(m)$ denote the similarity class in Λ_i to which m belongs, and $Z(\lambda)$ the set of terminal nodes which are descendant of at least one $m \in \lambda$.

A valuation $v_i : \Lambda_i \rightarrow \mathfrak{R}$ of player i is a mapping from M to \mathfrak{R} where $v_i(\lambda)$ refers to player i 's valuation of the similarity class $\lambda \in \Lambda_i$. A strategy σ_i for player i specifies for each $n \in N_i$, a distribution $\sigma_i(n)$ over $\Delta M_i(n)$. For any strategy profile σ , and $z \in Z$, I let $P^\sigma(z)$ denote the probability that z is reached given σ .

A valuation equilibrium requires that at every node $n \in N_i$, player i chooses a move belonging to a reachable similarity class with maximum valuation and that valuations are consistent with the strategy profile. I give here the definition without trembles and like other previously defined concepts, it can be refined adding trembling-hand perturbations (see Jehiel and Samet (2007) for details).

Definition 10 *A valuation equilibrium is a strategy profile σ such that*

- 1) *For all $n \in N_i$, $\text{Supp } \sigma_i(n) \subseteq \arg \max_{m \in M_i(n)} v_i(\lambda(m))$ and*
- 2) *For all $\lambda \in \Lambda_i$ such that $P^\sigma(Z(\lambda)) > 0$, $v_i(\lambda) = \sum_{z \in Z(\lambda)} P^\sigma(z) u_i(z) / P^\sigma(Z(\lambda))$.*

The interpretation of a Valuation Equilibrium is as follows. For given valuations, the criterion used in part 1 of Definition 10 requires (for player i) identifying the value of a move with the valuation attached to the similarity class to which the move belongs. I would say this is a fairly natural criterion, very much in line with how players would choose a move in a game like checkers or chess based on some coarse evaluation of the board position (such as the number of pieces ahead or a weighted version of it). Concerning part 2 of Definition 10, this is a consistency requirement similar to that introduced for ABEE. It can be thought of as follows. Assume the play has stabilized to some strategy profile σ . A statistician in charge of assessing the value for player i of the similarity class $\lambda \in \Lambda_i$ can, for each play, check whether a move in λ was chosen and in such a case add the obtained player i 's payoff in a box labelled λ . By considering the long-run average

payoffs put in box λ , the statistician will end up with the valuation $v_i(\lambda)$ as expressed in part 2 of Definition 10 whenever $P^\sigma(Z(\lambda)) > 0$.

The learning environment as just described is one in which players do not form beliefs about others' strategy, and instead try to assess the values to them of different courses of behaviors based on the performance they obtain in relation to these. While this feature is shared by models of reinforcement learning, a specificity here is that what players try to assess does not correspond to strategies as usually considered but to similarity classes. From another perspective, the idea of valuation as defined over similarity classes (and not strategies) should be thought of as a familiar object from the viewpoint of game playing programmes (see Samuel (1959) for first developing such programmes in checkers) in which nodes (or board positions) are typically assessed according to some simplifying device. From that perspective the novel feature here is the derivation of the valuations of similarity classes from the strategy profile through the consistency mapping as expressed in part 2 of Definition 10.

In Jehiel and Samet (2007), we prove the existence of Valuation equilibrium in general finite environments, we establish the link of the trembling-hand version of Valuation equilibrium to Subgame Perfect Nash Equilibrium when players use the finest similarity partition (see Jehiel and Samet (2005) for the convergence analysis of the corresponding learning model), and we illustrate through various examples how, for a number of similarity partitions, the Valuation equilibrium offers predictions that differ from those obtained with Nash equilibrium when varying the information structure.

How does the Valuation equilibrium (VE) relate to ABEE? VE was conceived as the reinforcement learning counterpart of ABEE. The focus on one own's performance in VE as opposed to others' behavior in ABEE has led to a different notion of coarseness in the two approaches, which is parameterized by the similarity partition in VE and the analogy partition in ABEE. Beyond that general (and somewhat informal) connection, I note that in some simple situations, VE can be viewed as an ABEE of an appropriately defined setting. This is in particular so when nature first decides which decision problem x arises according to some probability distribution $p(\cdot)$, and then in each problem x the same decision maker has several alternatives $a \in A(x)$ resulting in terminal node z with payoff

$u(z) = f(x, a)$.⁸⁷ In this case, the VE for any partitioning of the moves a at the various x can be captured as an ABEE of an appropriately defined setting. To see this, add a player after each (x, a) in charge of choosing the payoff $f(x, a)$, and consider the analogy partition such that if move a at x and move a' at x' were originally in the same similarity class, then the nodes xa and $x'a'$ at which the newly introduced player must choose the decision maker's payoff are in the same analogy class. It is readily verified that the VE of the original setting are isomorphic to the ABEE in the modified two-player game. A similar construction would not work for more general extensive-form games such as the stopping games discussed in Jehiel and Samet (2007) in which the same player must make several decisions in sequence.

Sampling equilibrium

Osborne and Rubinstein (1998) have introduced a notion of equilibrium for static games, that will be referred to as $S(k)$ or sampling equilibrium, in which players base their choice of strategy on the average performance obtained for each possible action in samples of size k .⁸⁸ By considering heuristics based on performances (as opposed to theories about the other players), their approach can be related to the tradition of reinforcement learning, and their steady state formulation of the $S(k)$ equilibrium can be interpreted as a limiting outcome of such a learning model.

But, a distinctive feature of their approach is that each individual player is viewed as basing his strategy on observations of finite samples (and the steady state is reached only at the population level). In deterministic situations (requiring both that strategies are deterministic and that the payoffs attached to the various action profiles are deterministic), the heuristic based on finite samples would lead to no mistakes. Also, if the sample sizes were infinite, there would be no mistakes either, even in non-deterministic environments, as a consequence of the law of large number. Mistakes arise for finite sample sizes and non-deterministic environments.

⁸⁷Such situations cover examples 1 and 2 in Jehiel and Samet (2007) (that are used among other things to illustrate the difference of VE with Sequential Equilibrium) as well as the experimental setting studied in Jehiel and Singh (2021).

⁸⁸See Osborne and Rubinstein (2003) for a belief learning counterpart of the $S(k)$ equilibrium which they refer to as sampling equilibrium and Spiegler (2006ab), Bianchi and Jehiel (2015)-(2020), Cherry and Salant (2020) for various applications and/or extensions.

To give an illustration of their concept, consider a symmetric 2×2 game with actions $\{x, y\}$ whose payoffs for player 1 are indicated for each action profile in the following table:

	x	y
x	a	c
y	b	d

Assume that all a, b, c, d are distinct and that $b > a$ and $c > d$ so that there is no symmetric pure Nash equilibrium. Consider a symmetric $S(1)$ equilibrium in which players make their choice of actions on the basis of a (payoff) sample of size 1 for each action. Formally, it specifies a probability p that x is played. There are four possible sample profiles of size 1 (for actions x and y) where each sample profile gives rise to a specific choice of action.⁸⁹ Steady state requires that the resulting probability that action x is chosen coincides with p . For example, when $c > b > d > a$, the symmetric $S(1)$ equilibrium satisfies:

$$p = 1 - p + p(1 - p)$$

or $p = \frac{-1+\sqrt{5}}{2}$, which typically differs from the mixed Nash equilibrium probability.

The sampling equilibrium can be viewed as formalizing the bias that agents tend to extrapolate from small samples to form their expectations (here about how good the various available actions are). The mistake in ABEE can also be related to a form of (erroneous) extrapolation, to the extent that players wrongly believe that the aggregate observation in an analogy class applies equally to all elements of the analogy class. While the extrapolation in the sampling equilibrium can be viewed as attributing the individual (sampled) observation to the general case, in ABEE it can be viewed as attributing the aggregate observation to the individual cases. Also, in ABEE, the observation is not based on small samples unlike in the sampling equilibrium. Clearly, one could combine the ABEE approach with the idea that in each analogy class, players observe finite samples and extrapolate from those, as in the sampling equilibrium.

⁸⁹With probability p^2 , x is perceived to give a and y is perceived to give b , leading to a choice of y given that $b > a$. With probability $(1 - p)^2$, x is perceived to give c and y is perceived to give d , leading to a choice of x given that $c > d$. With probability $p(1 - p)$, x is perceived to give a and y is perceived to give d , leading to a choice of x whenever $a > d$ and y otherwise. With probability $(1 - p)p$, x is perceived to give c and y is perceived to give b , leading to a choice of x whenever $c > b$ and y otherwise.

It may be worth highlighting that in both the sampling equilibrium and ABEE, the heuristic used by the agents can be viewed as being exclusively based on what these agents are assumed to have learnt about their environment: the aggregate distribution of behaviors in each analogy class in ABEE, and the finite sample for each action in $S(k)$. While in ABEE, the heuristic can be phrased as requiring that agents assume the general aggregate knowledge in a group of contingency applies separately to each contingency in the group, the heuristic in $S(k)$ can be phrased as requiring that the average payoff obtained in the sample is viewed as representative of how good the action is. This direct explicit mapping from the feedback to the subjective theories held by agents distinguishes ABEE and the sampling equilibrium from the Berk-Nash equilibrium but also from the self-confirming equilibrium in which such a mapping is not made explicit (see the above discussion of self-confirming equilibrium).⁹⁰

Limited foresight equilibrium

In the context of repeated alternate move games, Jehiel (1995) has introduced the concept of limited foresight equilibrium in which players are assumed to make forecasts only within a limited horizon. Here, the sophistication of a player is parameterized by how far ahead the player makes predictions. Given the stationarity of the environment considered in Jehiel (1995), players are assumed to consider the discounted sum of payoffs within their horizon of forecast as their criterion, and, in equilibrium, forecasts are assumed to be correct. Jehiel (2001) looked at repeated simultaneous move games, and considered stochastic criteria as an alternative modeling of how players assess the payoffs in the periods lying beyond the horizon of forecast.

To give a flavor of the concept, consider a two-player setting in which in every period t , the stage game payoff of player i is $u_i(a^t)$ where $a^t = (a_1^t, a_2^t) \in A = A_1 \times A_2$ is the action profile at t , and players discount future payoffs with a common discount factor δ . In alternate move games, player 1 can change his action a_1 in every odd period, player 2 can change his action in every even period, and actions remain the same until they are

⁹⁰The Valuation Equilibrium and the limited foresight equilibrium to be described next also share this property. Regarding the Bayesian Network Equilibrium, if one interprets it as resulting from the subjective causality relations considered by the decision maker, BNE does not share this property, but Spiegel (2017) offers another interpretation (with missing data) that would allow to include BNE as well. Of course, one could also use the ABEE interpretation of BNE and its associated feedback structure to do this as well.

changed. In standard repeated games, players can change their action in every period.

Player i is assumed to have a forecast horizon n_i . Player i decides his period t action a_i^t at history h on the basis of his forecast $f_i(a_i | h)$ about what will happen within his horizon of forecast if he plays a_i . In general, $f_i(a_i | h)$ is a distribution over action profiles from the current period t till period $t + n_i$ and player i 's forecast includes player i 's own actions in all future periods within the horizon of forecast.⁹¹ I let $v_i(a_i | f_i, h)$ be the criterion used by player i to compare the various actions a_i based on the forecasting rule f_i and the history h . In Jehiel (1995), I considered the deterministic criterion consisting of the discounted sum of payoffs within the horizon of forecast. In Jehiel (2001), I considered a stochastic version of this allowing the unpredicted time periods to contribute in a noisy fashion to the criterion. For the sake of simplicity, I will have in mind the deterministic criterion in what follows.

As usual, a strategy σ_i for player i specifies for every history h (where i must move), a probability distribution over the action space A_i assumed to be finite. I let $\sigma_i^h[a_i]$ denote the corresponding probability that a_i is played, and $\sigma = (\sigma_1, \sigma_2)$ denote a strategy profile. I also let $q(a_i | h, \sigma)$ denote the distribution over paths from h onwards generated by σ when player i plays a_i at h , and $[q(a_i | h, \sigma)]_{n_i}$ denote the marginal of $q(a_i | h, \sigma)$ over the next n_i periods after h . A limited foresight equilibrium is defined as follows:

Definition 11 *A strategy profile σ is a (n_1, n_2) -limited foresight equilibrium if for $i = 1, 2$ there exists a forecasting rule f_i such that for all h, a_i, a_i^**

- 1) *If $\sigma_i^h[a_i^*] > 0$ then $a_i^* \in \arg \max_{a_i} v_i(a_i | f_i, h)$ and*
- 2) *$f_i(a_i | h) = [q(a_i | h, \sigma)]_{n_i}$.*

While Condition 1 ensures that the strategy is related to the forecasting rule, as dictated by the criterion v_i , Condition 2 imposes that the limited forecasts are correct on and off the equilibrium path. In Jehiel (1995), I consider repeated alternate move games and further impose that the strategy of player i depends at most on the last N_i

⁹¹An alternative would be to limit the forecast to the reaction function of player $-i$, and allow player i to choose an optimal plan within the horizon of forecast. Given the rolling nature of the horizon of forecast, this alternative modelling would lead to time inconsistencies, which I find undesirable. Besides, I would argue that, in a number of contexts, the real object of choice at a given time period is the current action and not the plan of action, leading me to prefer the formulation of the forecast as a function of the current action (see Jehiel (1995) and Rubinstein (1998) for further discussion of this).

periods in addition to the calendar time. I show that limited foresight equilibria in pure strategies always exist, that they are cyclical and that strategies can only depend on the last action in addition to the calendar time. I also note that players may sometimes be better off with shorter horizons of forecast. In Jehiel (1998), I consider a learning model and establish the convergence of it to a limited foresight equilibrium, thereby providing a learning justification to the correctness of the forecasts assumed in the limited foresight equilibrium.⁹² In Jehiel (2001), I consider repeated simultaneous move games with noisy criteria, provide a general method to construct all limited foresight equilibria (this requires elaborations on the standard dynamic programming techniques), and observe that limited foresight may make cooperation easier to sustain than universal defection in repeated prisoner's dilemmas.⁹³

The limited foresight equilibrium like the other solution concepts discussed so far leads (sometimes) players to make suboptimal decisions, given that the criterion v_i in the limited foresight equilibrium need not always coincide with the rational criterion. Obviously, the mistakes in the limited foresight equilibrium are different from those in ABEE, but a common feature is that the criterion used by the players is related in a simple way to the object the players are supposed to learn (the analogy-based expectations in ABEE and the limited forecasts in the limited foresight equilibrium). It may be worth mentioning that the desiderata to have a simple formulation of the criterion v_i to be based solely on the forecasting rule f_i has led me to restrict the study of limited foresight equilibria to repeated games or repeated alternate move games for which, due to the stationarity of the environment, it was somehow natural to view the payoff obtained within the horizon of foresight as representative of the overall payoff. Moving away from such stationary environments, limited foresight equilibrium would require defining a criterion to assess the frontier nodes at the boundary of the forecast horizon. Such criteria are commonly used in game playing programmes, and the methodology developed for the valuation equilibrium could be combined with the methodology of the limited foresight equilibrium to cover such extensions. Of course, one could also relax the assumption that the forecasts are fully correct within the horizon of forecast, for example using the methodology of the

⁹²See Mengel (2014) for another learning model with limited foresight in finite horizon interactions.

⁹³Rampal (2017) or Ke (2019) develop alternative approaches to limited foresight which are based on introspection rather than learning.

analogy-based expectation equilibrium. The first unpublished version of ABEE allowed for such a combination (see Jehiel (1999)).

Behavioral equilibrium

In the context of adverse selection models of the "market for lemons" type, Esponda (2008) defines a solution concept that mixes the idea that the buyer fails to appreciate the link between the information of the seller and his behavior, as in the fully cursed equilibrium or the ABEE with coarse analogy partition, with the idea that the buyer forms his belief about the distribution of the quality of the good based on the payoff observation when there is a transaction. The naive behavioral equilibrium assumes that the buyer views the qualities observed when there is transaction as representative of the overall quality distribution, and adjusts his behavior optimally given this conjecture.

To give a concrete illustration of the concept, let me use the leading example in Esponda (2008).⁹⁴ The seller S is privately informed of the quality q of his good where q is assumed to be uniformly distributed on $[0, 1]$. The values of the good to the seller and the buyer B are $v_s(q) = q$ and $v_b(q) = q + b$ (with $0 < b < \frac{1}{2}$), respectively. The seller and buyer simultaneously quote a price p_s , for the seller, and p_b , for the buyer. There is transaction whenever $p_b > p_s$, and the transaction price is p_b so that quoting price $p_s = v_s(q)$ is a weakly dominant strategy for the seller.

The price p^{NBE} quoted by the buyer in a naive behavioral equilibrium satisfies:

$$p^{NBE} = \arg \max_{p_b} [E(q + b \mid q < p^{NBE}) - p_b] \Pr(q < p_b)$$

where the term $E(q + b \mid q < p^{NBE})$ represents the experienced value of the good for the buyer when there is transaction (i.e., when $q < p^{NBE}$) and $\Pr(q < p_b)$ represents the probability of getting the good when quoting p_b that is correctly computed on the basis of the observed distribution of p_s . Solving for p^{NBE} yields $p^{NBE} = \frac{2b}{3}$, which is less than b (the Nash equilibrium price p_b) and a fortiori (when $b < \frac{1}{2}$) less than $\frac{2b+1}{4} = \arg \max_{p_b} [E(q + b) - p_b] \Pr(q < p_b)$ the fully cursed equilibrium price p_b . Unlike in the fully cursed equilibrium, the buyer is not mistaken about the equilibrium value of

⁹⁴The naive behavioral equilibrium as defined more generally in Esponda (2008) need not always exist, thereby justifying the focus on the lemon setup where it is well defined (see Esponda (2008) for further elaborations on this).

the good when he buys it (since the term $E(q + b \mid q < p^{NBE})$ correctly represents this value), and the mistake comes from the buyer's failure to appreciate that when he changes his quoted price p_b this affects the average quality of the transacted goods. This failure leads to less trade than in the rational case (as downward -resp. upward- deviations are perceived as weakly more -resp. less- profitable than they really are).⁹⁵

How does the naive behavioral equilibrium relate to ABEE? A priori, in the naive behavioral equilibrium, coarse agents mix two kinds of feedback, one on the distribution of others' actions and the other on their own payoff, which does not seem compatible with the machinery of ABEE (this mixture is also responsible for the inexistence of naive behavioral equilibrium in some cases). Yet, in the context of the adverse selection problem just outlined, it is possible by a suitable choice of extensive form game and analogy partitions to view the NBE as an ABEE.

To see this, consider the following extensive form game. The sellers' type is q as usually defined. Then players S and B simultaneously choose p_s and p_b , as described above. When there is transaction, i.e. $p_b > p_s$, another player called N_q chooses the quality or equivalently the value to the buyer specified to be $v_b^n = v_b(q) = q + b$ where q is the type of the seller. And the utility of the buyer is defined to be $v_b^n - p_b$ when there is transaction and 0 otherwise. When all the decision nodes of player N_q form a single analogy class, and all the decision nodes of the seller form another analogy class, it is readily verified that the ABEE so obtained coincides with the naive behavioral equilibrium as previously defined. From a broader perspective, the ABEE in its extensive form variant allows the modeler to capture effects of the selection bias type as arising in the naive behavioral equilibrium or the investment problem studied in Jehiel (2018).

The equilibrium approaches to be described next are somehow less closely related to ABEE. I include a discussion of them in an attempt to offer a broader perspective on how complexity/bounded rationality issues have been approached in the literature, while each time indicating the difference of perspective with ABEE.

Equilibrium with strategy restrictions

When the strategic environment is too complex, it sounds plausible that players would

⁹⁵If one observes more trade than the rational approach predicts, one cannot rely on the naive behavioral equilibrium to explain it (see Fudenberg and Peysakhovich (2016) for an experiment on this).

only consider a restricted set of strategies, and choose the best strategy within this restricted set. It is then natural to consider the Nash equilibria of the game appropriately modified to account for the restrictions in the strategy sets.

Formally, consider a strategic environment with players $i = 1, 2$. Denote by Σ_i the set of all possible strategies of player i , and by σ_i a representative strategy in Σ_i . Let $u_i(\sigma_i, \sigma_{-i})$ denote the expected utility obtained by player i when player i plays according to σ_i and player $-i$ plays according to σ_{-i} . Observe that such a formalism may accommodate applications in which player i would engage in several different interactions in which case a strategy would allow the player to adjust his behavior to the interaction, or interactions with private information in which case a strategy would allow the player to adjust his behavior to his private information, or multi-stage interactions (with private information) in which case the strategy would allow the player to vary his behavior as a function of history (and the private information). In this reduced-form presentation, a strategy can be thought of as a complete (possibly contingent) plan of actions to be chosen at an ex-ante stage before the plan is executed.

Restrictions in strategy sets can be modelled by assuming that player i only considers strategies σ_i in the subset $\Sigma_i^s \subseteq \Sigma_i$. An equilibrium with such strategy restrictions is a strategy profile $\sigma^s = (\sigma_i^s, \sigma_{-i}^s) \in \Sigma^s$ such that for each i , $\sigma_i^s \in \arg \max_{\sigma_i \in \Sigma_i^s} u_i(\sigma_i, \sigma_{-i}^s)$. That is, it is a Nash equilibrium of the normal form game defined by the utility functions u_i, u_{-i} and the action spaces Σ_i^s and Σ_{-i}^s , for players i and $-i$, respectively. As such, it is not a novel solution concept,⁹⁶ and the challenge is to propose reasonable choices of Σ_i^s and Σ_{-i}^s in each application as well as to understand how such strategy restrictions affect the equilibrium analysis. Equilibrium with strategy restrictions is the route followed by Compte and Postlewaite (2019) to revisit a number of classic economic applications.

The potential appeal of the approach from a bounded rationality perspective is that players have fewer aspects to learn than in the usual Nash equilibrium case. Indeed, assuming the behavior of player $-i$ does not change (as should be the case in steady state), it is enough for each player i to find out the optimal strategy in the restricted set Σ_i^s as opposed to the entire strategy set Σ_i (thus suggesting an easier maximization/learning process). I made a related comment for ABEE when I noted that in ABEE, player

⁹⁶Existence of Nash equilibrium with strategy restrictions is not guaranteed in general and it typically requires some conditions on the sets Σ_i^s (in particular, related to convexity).

i only has to learn the aggregate behavior in each analogy class, which is easier than having to learn the behavior in every contingency separately. But, the underlying learning environments are different in the two cases. While ABEE finds its motivation in (coarse) belief-based learning models in which learning occurs at the population level (see the above learning motivation for ABEE), the equilibrium with strategy restrictions is better viewed as the outcome of a reinforcement learning model in which the same players would be playing many times and these would seek to learn the performance of the various strategies in their consideration set.⁹⁷

A critical aspect in the equilibrium with strategy restrictions concerns the choice of the restricted strategy sets Σ_i^s . Compte and Postlewaite (2019) consider a variety of illustrations many of which amount to reducing the strategy sets to one-dimensional families of strategies that they think of as plausible restrictions for their applications. Despite these many illustrations, I think fair to say that Compte and Postlewaite do not provide a general recipe for how to choose such strategy restrictions, making in my view the use of their approach not straightforward.⁹⁸

It should also be mentioned that the equilibrium with strategy restrictions has precedents in the literature. For example, in the context of multi-stage games, Neyman (1985) considers restricting the players to use strategies that can be described by automata using at most n states. With a similar machinery but a different context in which players would be exposed to different interactions, one may consider that players categorize the interactions (or the various branches of those) into a finite number of categories while imposing that the behavior is the same within the same category (using an approach similar to

⁹⁷There has been a long debate in the learning literature whether belief-based or reinforcement learning models explain better observed behaviors (see in particular, Camerer and Ho (1999)). I would suggest the answer to this question depends on the environment one has in mind. If the same player keeps playing again and again, and if he has easily access to his performance, then reinforcement learning models are likely to be more appropriate. On the other hand, if players do not play themselves many times but are instead exposed to some statistics describing past play (as for the motivation of ABEE), reinforcement learning models seem inappropriate. I would add that even if the same player keeps playing again and again, his performance may be observed with a lag and in a noisy fashion, making it less natural that players would rely on reinforcement learning models in such contexts.

⁹⁸As already mentioned, I think in the ABEE context, the choice of analogy partitions is less subject to a similar critique whenever the analogy partitions are viewed as representing exogenously coarse feedback. See also the next section for elaborations on the choice of analogy partitions when these are viewed as learning facilitating devices.

that of Rubinstein (1986),⁹⁹ Samuelson (2001) and Mengel (2012) develop such models, see also the rule of thumb theme developed by Rosenthal (1993) and the rule rationality theme developed by Aumann (2008) for related ideas).

Equilibria with imperfect optimization

Another approach to the complexity of finding best-responses is to allow for mistakes in the best-responses. One such approach is the ε -equilibrium in which a player is only required to choose a strategy that induces an expected payoff no more than ε away from the maximum achievable payoff (see Radner (1980) who uses this concept to explain cooperation in finite horizon interactions). The ε -equilibrium can be related to the satisficing idea of Simon (1956) according to which a decision maker is viewed as being satisfied with his strategy, as long as it delivers an expected payoff no less than his aspiration level.¹⁰⁰ In the ε -equilibrium, the aspiration level of a player is set ε below the maximum payoff this player can achieve in equilibrium. While it is not clear, why the aspiration level would be set in this way, it should be noted that if it were set above the maximum achievable payoff, the considered strategy profile would not be in equilibrium (given that the player would then not be satisfied with his strategy). The ε -equilibrium can then be viewed as offering a simple way to capture satisficing behaviors in steady state environments (even if with this interpretation, there is no reason to assume that ε is small, as usually considered in the literature).

Another approach to imperfect optimization consists in allowing players to stochastically choose each of their various available actions while requiring that those actions involving larger payoff losses would be chosen with lower probabilities. This is the idea behind the quantal response equilibrium (QRE) developed by McKelvey and Palfrey (1995)). The most popular parameterization of QRE follows the logit specification. In a normal form game with finite action space A_i and utility function $u_i(a_i, a_{-i})$ for player i , σ is a

⁹⁹Rubinstein (1986) considers a setup in which players care about the complexity of their strategy in addition to their payoff as usually defined where the complexity is measured by the number of states that the automaton representation of the strategy requires. It should be mentioned though that the lexicographic criterion considered by Rubinstein (1986) does not allow to view his concept as an equilibrium with exogenous strategy restrictions. Rubinstein's concept can be viewed as a Nash equilibrium in which one would have to modify players' utility to include players' concern about the simplicity of their strategy.

¹⁰⁰See Gigerenzer (2001) for more recent recent developments of the satisficing approach, this time applied to heuristics that can be viewed as focal strategies on which players focus their attention (similarly as in the restricted strategy set approach discussed above).

QRE with noise parameters $\lambda = (\lambda_i, \lambda_{-i})$ if the probability that a_i^* is chosen satisfies

$$\sigma_i[a_i^*] = \frac{\exp \lambda_i V_i(a_i^*; \sigma_{-i})}{\sum_{a_i} \exp \lambda_i V_i(a_i; \sigma_{-i})}$$

where $V_i(a_i; \sigma_{-i}) = \sum_{a_{-i}} \sigma_{-i}[a_{-i}] u_i(a_i, a_{-i})$ is the expected utility of player i when he plays a_i and player $-i$ follows σ_{-i} . When $\lambda_i = 0$, player i randomizes uniformly across all actions, and when $\lambda_i = \infty$, player i plays an exact best-response to his environment. When all players i use $\lambda_i = \infty$, the QRE is a Nash equilibrium. Intermediate values of λ_i accommodate situations in which player i is responsive to payoff differences but in a smooth way.

QRE can be viewed as formulating in a game theoretic context, familiar ideas from the discrete choice literature (McFadden (1973)). From that perspective, the criterion used by players may be thought of as a reduced form representation where the true preference of player i attached to action a_i would be the sum of $V_i(a_i; \sigma_{-i})$ and some idiosyncratic term. As shown by McFadden, when the idiosyncratic term is independent across actions and follows an extreme value distribution, one obtains the logit formulation. Viewing the QRE in this way makes this concept in fact coincide with the Bayes Nash equilibrium of the game in which the idiosyncratic terms would be private information to the player. Some researchers interpret QRE differently, viewing the noisy criterion as a manifestation of mistakes (see in particular Chen et al. (1997)). But, this interpretation requires that players would be able to adjust the magnitude of the mistake to the payoff loss, and it is not clear through which mechanism boundedly rational agents could achieve this.

The ε -equilibrium as well as the QRE are clearly touching on aspects of bounded rationality that differ from those considered in ABEE to the extent that they do not challenge the rational expectation assumption.¹⁰¹ In my view, the QRE methodology is quite useful to analyze experimental data which are inevitably noisy. As such, it can

¹⁰¹See however an alternative to QRE proposed by Friedman and Mezzetti (2005) and more recently by Friedman (2019) who consider putting (white) noise in the belief about σ_{-i} rather than in the criterion (see also the M -equilibrium introduced by Goeree and Louis (2018) which combines both noisy best-responses and partially erroneous beliefs). Note that except for the sampling heuristic discussed above, these papers (in particular Friedman (2019)) do not provide explicit heuristic mechanisms through which the distribution of noisy beliefs would be centered around the correct distribution.

naturally be combined with other approaches including ABEE, as illustrated in Huck et al. (2011).¹⁰²

Mean field equilibrium

When there are many players such as in trading environments or in complex network environments, it is somehow natural for players to summarize the state of the economy by some moments of the distribution of plays where such moments as well as possibly individual characteristics are viewed as determining the individual preferences. This is the approach pursued in mean field game theory (see Lasry and Lyons (2006)) where it is assumed that individual players take the moments and their dynamics as given, optimize against these, and the evolution of the moments are endogenized by aggregating the individual behaviors, as in rational expectation models.

The ABEE shares with the mean field game theory the idea that players base their strategy on a simplified representation of their environment. But, a crucial difference is that while in the mean field game theory, the aim is to study rational agents and the simplification is viewed as a legitimate approximation of Nash equilibrium when there are many players, the aim of ABEE is to propose a framework to describe different forms of bounded rationality (as obtained by varying the analogy partitions) and study their implications for equilibrium predictions.

To give an illustration of the difference, consider a game with many players each having to make a binary decision $a \in \{0, 1\}$ such as a technology adoption in a network. In the tradition of mean field game theory, one could consider anonymous environments,¹⁰³ and assume that players care about the adoption rate of their neighbors as well as their number. One could then assume that players make their choice of technology assuming every neighbor follows the aggregate adoption rate, and that the number of neighbors is drawn from a distribution that respects the degree distribution as observed in the true network. The resulting equilibrium would approximate the Bayes Nash equilibrium of a setup with many players in which individuals would make no observation about the network assumed to be formed stochastically in agreement with the degree distribution.

One could instead assume that players observe very well the specific network they

¹⁰²Haile et al. (2008) challenge the predictive power of QRE if one allows for arbitrary distributions of the noise terms. Similar considerations can be developed when combining QRE with other approaches.

¹⁰³Anonymous environments are also the subject of the large games studied by Kalai (2004).

are in, or at least the number of neighbors their neighbors and themselves have, but to form their expectation about the adoption rate of their neighbors, that they rely on the aggregate adoption rate in the network. Now, the adoption rate of a player may depend on the number of his neighbors, since this number is observed, but players' expectations about their neighbors' adoption rates would generally be incorrect since every player would assume his neighbors all behave according to the aggregate adoption rate when in reality their adoption rate would be influenced by their degree. That kind of environment is well captured by ABEE assuming, for the given network, that players use the coarsest analogy partition that puts together all players in the network. In some sense, it is suggestive of how the mean field game theory agenda could be extended to cover situations in which the mean field approach would allow to approximate ABEE rather than Nash equilibria.

I have so far discussed equilibrium approaches. I end this Section with a discussion of two non-equilibrium approaches initially introduced in Artificial Intelligence and psychology that rely on learning procedures that share some common motivation with ABEE and/or related concepts. These are the Case-Based Reasoning and the Learning Classifier Systems.

Case-Based Reasoning

Case-Based Reasoning as introduced by Schank (1982) in AI is, in broad terms, the process of solving new problems based on the solutions of similar past problems. It requires memorizing past cases, identifying which past cases are similar to the problem of interest, and possibly adapting the solutions that worked well in past similar cases to the current problem. Relatedly, but with a simpler focus on decision theory as conceived by economists, Gilboa and Schmeidler (1995) have developed a theory in which in order to evaluate an action a^* in a new problem p , the decision-maker would look at past cases M consisting of triples (q, a, r) where q denotes a problem previously encountered, a denotes the chosen action in q and r denotes the outcome, and he would assess a^* in p according to the criterion:

$$U_{p,M}(a^*) = \sum_{(q,a^*,r) \in M} s(p,q)u(r) \quad (6)$$

where $s(p,q)$ is the similarity between problems p and q and $u(r)$ is the utility derived

from the outcome r . That is, the decision-maker assesses a^* in problem p according to the weighted sum of utilities he derived in the various past problems q when a^* was chosen where the weight given to past problem q is defined by the similarity $s(p, q)$ between q and p .

Case-Based Reasoning shares with the learning approach I used to motivate ABEE, the idea that observations made in past possibly different situations would be used to make assessments and decisions in the current situation of interest. In both Case-Based Decision Theory and ABEE, whether and to which degree other situations are used is given exogenously by the analogy partition in ABEE and the similarity function in CBDT.¹⁰⁴ But, there are important differences both in terms of the pursued methodology and in terms of objectives. First, ABEE is concerned with strategic environments and past cases are used to assess frequencies of opponents' actions over different analogy classes. As such, it builds on belief-based learning models whereas Case-Based Reasoning is closer to the reinforcement learning tradition in which past cases are used to assess how good the various possible options are. From that perspective, CBR shares more with Valuation Equilibrium than ABEE. Second, ABEE is concerned with steady state in which lots of data have been accumulated whereas CBR and CBDT are concerned with the resolution of new problems. This has led to an equilibrium formulation in ABEE in which the beliefs are related to the strategies and the analogy partitions through a consistency requirement (similarly as in Nash Equilibrium) when there is no such endogeneity in CBDT. Third, ABEE can be viewed as relying on a frequentist approach in which subjects are exposed to frequencies of actions in the various analogy classes whereas such a frequentist interpretation is less readily available for CBDT.¹⁰⁵ I would argue that the frequentist approach makes it easier to interpret ABEE as resulting from subjects being exposed to coarse statistics (possibly as given by society) whereas such an interpretation would be less natural

¹⁰⁴I note that the Case-Based Reasoning of Schank suggests the additional desiderata that only analogies that help achieve better outcomes should be considered and that past solutions that worked well in previous problems may need to be adapted to the current problem for efficiency purposes. I would interpret the latter idea of adaptation in relation to the idea that the notion of analogy should be extended to actions.

¹⁰⁵CBDT can accommodate the frequentist approach by imposing that the similarity is either 0 or 1 and by normalizing expression (6) by the number of similar cases in which action a^* was chosen. While this is discussed by Gilboa and Schmeidler, it is not considered as the main model of CBDT, which allows for general similarity functions.

in CBDT (in CBDT, M is what the decision maker remembers from the past and s is the subjective similarity function that he considers). Finally, and in relation to Schank's view, I note that the main objective of ABEE consists in describing better the actual behaviors of economic agents, which contrasts with the AI objective of finding good solutions to new problems. This has led me to use ABEE as a tool to explore whether observed puzzles could receive reasonable explanations and to not necessarily worry whether subjects could be performing suboptimally.

Learning Classifier Systems

Learning Classifier Systems (LCS) as introduced by Holland (1992) consider complex decision problems in which the system may evolve from state to state depending on the chosen actions and in each state the agent may consider various rules of thumb or classifiers to make his choice of action. LCS allow for new generations of classifiers, but for now I will assume that the classifiers are exogenously given (I will come back to the generation of new classifiers in the next Section). At every step t of the learning process, different strengths or indices are attached to the various classifiers. When at some state s , the agent is required to choose the available classifier with maximum strength, call it C . Then some instantaneous payoff observation u_t is made and the system moves to another state s' . The strength of the classifier C is then adjusted from $v_t(C)$ to $v_{t+1}(C)$ according to some rule in the spirit of dynamic programming. For example, in contexts with standard exponential discounting, a typical updating rule would require that $v_{t+1}(C)$ is a convex combination between $v_t(C)$ and $u_t + \delta v_t(C')$ where δ is the discount factor and C' is the classifier that would be chosen at s' on the basis of the profile of the time t strengths (as given by $v_t(\cdot)$).

The main objective of such learning systems as viewed by machine learning researchers is to find out or at least approximate optimal solutions. The Q-learning algorithm defined for Markovian processes in which the classifiers consist of the set of all possible actions at every state (similarly as in standard reinforcement learning models) leads asymptotically to the optimal solution when the adjustments of the classifier strengths is slow enough (see Watkins and Dayan (1992)). But, for other contexts and more limited sets of classifiers, such convergence results need not hold.

The main objective of researchers dealing with bounded rationality in economics is different and can be related to whether an approach like LCS can give rise to predictions different from the ones arising with rational expectations. To my knowledge, there have not been many such uses of LCS in economics with the notable exceptions of Marimon et al. (1990) and Lettau and Uhlig (1999) in macroeconomics. For the purpose of the present paper, let me now abstract away from convergence issues, and assume, for a given set of classifiers, that the classifier strengths have converged. In such a limiting situation, the strengths of the classifiers would satisfy a consistency requirement that relates those strengths to the true underlying process including the effective choices of classifiers at each state. Moreover, at every state, the agent would choose a classifier with maximal strength. The obtained limiting outcome is somehow related to the Valuation Equilibrium introduced above if one identifies a similarity class and its valuation in VE with a classifier and its strength in LCS. But, there are some differences between the two. First, the similarity class in VE is viewed as dictated by the kind of feedback (about past problems) provided to subjects whereas the classifiers are viewed as (available) local models. Second, I allowed for multiple players in VE which is not so in LCS (even though extensions of LCS to multiple players could be considered). Third, consistency in LCS would not be defined as I considered it for VE. In VE, I assumed that the total payoff obtained in the game could directly be related to the similarity class (thereby leading to consistency as shown in Definition 10). In LCS, consistency would relate the strengths of the various classifiers through the dynamic programming formula. While an exact and general formulation of consistency for LCS is to be developed, I note that the two approaches to consistency in VE and LCS would correspond to environments that differ whether total payoffs can be observed quickly (for VE) or not (for LCS).

5 What's next?

In this Section, I indicate what I regard as possible next steps of the ABEE research agenda. Of course, this list is in no way exhaustive, as there are surely many other relevant directions for future research.

5.1 How to choose analogy partitions?

I have indicated three different ways of interpreting the analogy partitions in subsection 2.4. When analogy partitions are viewed as the result of exogenous feedback limitations, there is no reason to seek for further endogenization except if these limitations can be influenced by some economic agents. In subsection 5.1.1, I suggest such a perspective. When analogy partitions are viewed as a way to formalize psychological biases, I would say the input of psychologists is probably the most useful to better understand what aspects of contingencies agents find most salient. Finally, when analogy partitions are viewed as a device to simplify learning, I would say more is needed to understand how the choice of analogy partitions may facilitate learning. In subsection 5.1.2, I suggest some reduced-form approaches to this.

5.1.1 Analogy partitions as influenced by third parties

When analogy partitions are viewed as resulting from the exogenously coarse feedback players have access to, sometimes some economic agents can have a control on this. If one thinks societies can exert such a control, then a natural question is how the feedback should be chosen to maximize welfare or some other measure of efficiency. This can be viewed as being related to the debate on transparency (an attempt to impose increasing criteria of transparency to economic actors such as banks, firms etc), but note that transparency in most of the economics literature is generally thought of in relation to the information transmission within the interaction of interest rather than feedback related to data concerning previous interactions as in the ABEE framework. If on the other hand, one thinks economic agents possibly with non-benevolent interest can exert such a control, it may suggest a new role for regulation, as there is no reason that the best-choice of analogy partitions from a private interest viewpoint would coincide with what society likes best (the study of manipulative auction design in Jehiel (2011) can be viewed as providing an illustration of this).

Moving one step further, one could also envision situations in which different agents would propose different aggregate statistics to players. A question then arises as to how players faced with competing aggregate statistics would behave. I will discuss possible ways to cope with this in the next subsection. Once this is settled, one would still have to

understand the equilibrium behavior of competing economic agents offering such aggregate statistics.

5.1.2 Analogy partitions as chosen by individuals

In this part, I assume that players exert some control on the choice of analogy partition that they consider.

Analogy partitions, statistics, and machine learning

When viewing the motivation for analogy partitions as simplifying devices to interpret rich datasets, an obvious branch of literature that sounds relevant is machine learning.

One fundamental idea in statistics and machine learning is the so called bias-variance trade-off.¹⁰⁶ If contingencies are categorized too finely, the obtained statistics in each category are not reliable enough, and if they are categorized too coarsely, the statistics may be less relevant for adjusting the best-response. One possible reduced-form approach to this in the ABEE setting would be to require that analogy classes be reached with a probability no smaller than some positive threshold where this threshold would be viewed as a primitive of the model. I have considered in the end part of Jehiel (2005) a related idea where I imposed in extensive form games that off-the-path contingencies should be bundled with some on-the-path contingencies. Going beyond this, when one can define a natural distance between contingencies, the bias-variance trade-off suggests that players would bundle contingencies according to how close they are as in the neighborhood analogy partitions mentioned in Section 2.5.1 with the extra constraint that each analogy class should be reached with a sufficiently high probability.¹⁰⁷

Another classic tool in machine learning is the so called K-means clustering, which addresses the question of how to best cluster n points into K categories so as to minimize the sum of the square distance between each original point and the representative element of the cluster to which the point is assigned. This method has been extended to cover other

¹⁰⁶In non-strategic contexts, Fryer and Jackson (2008) have used this trade-off to model discrimination. Mohlin (2014) has developed a model to best solve this trade-off in simple non-strategic settings. In relation to this trade-off, Olea et al. (2019) have suggested that models with fewer explanatory variables may subjectively be considered as more reliable in the face of small datasets.

¹⁰⁷Jehiel and Mohlin (2021) discuss a related idea in adverse selection environments involving cyclical behaviors.

notions of closeness such as the Kullback-Leibler divergence when points are distributions (see in particular Banerjee et al. (2005)). In Jehiel and Weber (2021), we use the clustering idea to propose an endogeneization of the analogy partitions in contexts involving n different stage games. In doing so, we identify a cluster in machine learning with an analogy class in ABEE. Like in the K means clustering, we postulate that players rely on K analogy classes, and the choice of analogy partition should satisfy the property that the distribution of opponent's behavior in a game γ is closest to the average distribution in the analogy class $\alpha(\gamma)$ to which γ is assigned than any other average distribution as appearing in the other analogy classes α' considered by the player. We refer to this property as a stability idea, as it ensures that a player checking from time to time whether a given game γ is well assigned given the considered analogy partition would stick to his original analogy partition. Compared to the classic clustering idea, there is an extra endogeneity which is related to the ABEE formulation. Each grouping or clustering affects the distribution of points, unlike in the classic machine learning exercise in which the points are exogenously given. This endogeneity in turn makes the existence of pure stable ABEE problematic for some situations, in which case one has to extend the definition to allow for mixed stable ABEE. We have applied the approach to the study of symmetric two-player games with one-dimensional action spaces, and derived qualitatively different insights according to whether players' actions are strategic substitute or complements.

One may also mention the machine learning ideas developed in genetic algorithm (see Holland (1975)) to generate new rules of thumb or classifiers from old ones as a possible route to induce new choices of analogy partitions. One possible idea I see in relation to this is that if the variance of behavior observed in a given analogy class is too large, the player could somehow split the analogy partition into several pieces, and if the aggregate behaviors observed in two analogy classes are close enough to each other, the player could decide to merge them. There are clearly other ideas that could be considered in relation to this.¹⁰⁸ Obviously, more work is needed to understand this better.

Optimization approaches

As an alternative to machine learning, one could adopt more rational approaches to the

¹⁰⁸In particular, the LCS literature suggests ideas in the spirit of Axelrod's tournament (see Axelrod (1984)).

choice of analogy partitions. One such approach is that of a two-stage procedure. First, players choose their analogy partitions, then players play the corresponding ABEE. A question arises as to how players assess the effect of their choice of analogy partitions in the first stage, and how ABEE is defined in the second stage when a player considers arbitrary (including off-the-equilibrium path) analogy partitions. One option that I discuss at the end of Jehiel (2005) is that the ABEE played in the second stage corresponds to the choices of analogy partitions in the first stage (whether on or off the equilibrium path) and that the criterion for the first stage choice is the true expected payoff (it also corresponds to the approach in Heller and Winter (2020) who use this methodology to endogenize general families of biased beliefs in normal form games).¹⁰⁹ Several concerns can be raised about this approach. First, I find the commitment feature resulting from the assumption that the ABEE is the one corresponding to the individual choice of analogy partitions problematic in a number of applications, as the aggregate feedback in the ABEE is more naturally interpreted as coming from many different players (not just oneself).¹¹⁰ Second, it is not clear how players would have the correct understanding about how their choices of analogy partitions translate into true payoffs.¹¹¹ One could invoke some evolutionary or imitation process to give some rationale to this, but this would require, for evolution, that the evolutionary pressure operates at the level of analogy partitions, or, for imitation, that others' choices of analogy partitions can be observed in addition to their performances. While this is not impossible, I think fair to say that these conditions are not obviously met.

A somewhat different route is related to the rational inattention approach. One interpretation of the coarse feedback is that the agent does not pay full attention to every detail in the feedback (assumed to take the form of exhaustive and unprocessed raw data

¹⁰⁹I used this approach to obtain that in some cases, players would not adopt the finest analogy partition even if possible because it would give them some commitment power to choose a coarser analogy partition that turns out to be favorable in some interactions. This is similar to observations made in Dekel et al. (2007).

¹¹⁰One can drop this commitment aspect but if one places no restrictions on the analogy partitions, then it becomes weakly dominant to choose the finest analogy partition. Of course, it may be more reasonable to put constraints on the analogy partitions that players can consider, but then the equilibrium choice of analogy partitions will heavily hinge on these constraints.

¹¹¹This is somehow related to the infinite regress problem discussed in Mongin and Walliser (1988), Lipman (1991) and Conlinsk (1996).

about past behaviors and other aspects of the interaction) and focus on some aspects of it.¹¹² In the tradition of rational inattention models (à la Sims (2003)), one could attach a cost to processing more finely the information conveyed by the feedback, and one could assume that players choose their analogy partitions optimally taking into account the processing cost associated to the chosen analogy partition. I personally find the rational inattention approach to the choice of analogy partitions very demanding in terms of the level of rationality imposed on the players. It may be that the rational inattention approach can be viewed as a reduced form for a cognitively less demanding process, but a full-fledged argument supporting this view is still to be developed.¹¹³

Choosing among competing analogy partitions

In many instances, one may not have the full freedom to choose from the entire set of all analogy partitions, and one may only be exposed to the feedback corresponding to just a few analogy partitions, for example as provided by different experts or different statistical institutes. One possibility for the agent is to try to combine the information conveyed by the various types of feedback. Such a combination could take the form of attaching different weights (or probabilities) to the various analogy partitions, and assess strategies according to the perceived expected utility that results from this. Alternatively, the agent may only consider the most likely analogy partition, and optimize only the corresponding subjective utility. Note that in multi-stage settings, this selected analogy partition may change as more observations are made in the interaction.¹¹⁴ As yet another alternative,

¹¹²The salience perspective suggested above is related to this but with a psychological rather than rational viewpoint.

¹¹³The sparsemax model developed by Gabaix (2014) in the spirit of LASSO suggests a simplified version of the rational inattention model in which utilities are approximated by their second order Taylor expansion in models with continuous variables. While the simplification alleviates some of my concerns related to the rational inattention approach, it is unclear to me whether the proposed Taylor approximation can be viewed as capturing a behaviorally sensible heuristic.

From another perspective, Gagnon-Bartsch et al. (2020) develop an approach that they refer to as channeled attention that can be viewed as similar to the rational inattention approach except that it is based on a possibly misspecified prior belief of the agent. Such an approach can explain why agents would pay limited attention to the data and as a result be unable to detect the possible misspecifications of their prior. But, as with models of subjective priors or models with misspecifications (such as the Berk-Nash equilibrium), one would need to have a better understanding of how the misspecifications are formed in applications to make good use of such an approach.

¹¹⁴Mullainathan et al (2008) adopt the maximum likelihood approach to study a model of persuasion with coarse reasoning of the ABEE type. See also Schwartzstein (2014).

agents could choose the analogy partition that subjectively promises the highest expected payoffs (where the computation for a given analogy partition adopts the subjective representation attached to this analogy partition). This is similar in spirit to the optimal expectation framework of Brunnermeier and Parker (2005) with the constraints that beliefs should correspond to those generated by the available analogy partitions.¹¹⁵ In my view, more work including on the experimental side is needed to better understand how agents make use of different sources of feedback.

5.2 Learning to play ABEE

An important question not addressed so far is when learning dynamics of the type suggested after the definition of ABEE would converge to a steady state (and if so which one).¹¹⁶ This is a similar question to that pursued in the theory of learning in games (see, for example, the textbooks of Fudenberg and Levine (1998), Weibull (1997) or Young (2004)) with a focus on ABEE rather than Nash Equilibrium. In the standard case, dominance solvable games are known to induce nice learning convergence. The setup in Huck et al. (2011) suggests a similar notion for ABEE. In the example considered there, if one iteratively eliminates those strategies that are strongly dominated against all possible beliefs that are consistent in the ABEE sense for some remaining strategy profile, one eventually pins down a single strategy profile. This very property guarantees in turn the convergence of most learning models to ABEE, similarly as in the standard case. More work should be done to develop convergence results in other classes of games.¹¹⁷

I end this Subsection with two classic observations made in the context of learning in games, and I discuss how they could affect ABEE. First, some of the learning literature has suggested that the correlated equilibrium rather than the Nash equilibrium should be the right benchmark in the classic learning scenario with a single normal form game (see for example, Hart and Mas-Colell (2003) or Hart (2005)). It should be stressed that such

¹¹⁵Using the Bayesian network machinery, Eliaz and Spiegler (2020) develop such an approach to analyze a model of competing narratives.

¹¹⁶If these do not converge, then dynamic studies of the type considered in agent-based models may be more appropriate to describe long-run behaviors.

¹¹⁷Another class of games for which belief-based learning models of the fictitious play type is known to have good convergence properties is the class of 2x2 games (see Miyasawa, 1961). Li Calzi (1995) provides a similar result for families of 2x2 games in which players would use the coarsest analogy partition.

conclusions concern learning dynamics in which players learn as they keep playing against the same opponents (and the common history can then play the role of a coordinating device used to correlate the strategies of the players). In learning dynamics in which a given subject assigned to the role of a player plays just once as considered to motivate the ABEE, there is no common history on which players can coordinate their choice of strategy, and learning operates only at the level of the population. In this case, the ABEE formulation seems to be the appropriate long-run notion (exactly as Nash equilibrium would be in the standard case). Of course, one could well consider scenarios in which there is correlation between the types of the various players in the matching process, but this is covered in the ABEE formulation to the extent that any joint distribution of types (concerning both the regular and the cognitive dimension of the type) is allowed.

Second, a key insight from the literature on learning in multi-stage games is that unless the trembles are viewed as arising exogenously, it may not be obvious to justify the correctness of beliefs off-the-path, as assumed in traditional concepts such as Sequential Equilibrium. The same comment applies to ABEE, but it should be noted that the coarser the analogy partition, the less likely it is that an analogy class would be unreached in equilibrium. Such an observation can be used to restrict attention to analogy partitions and ABEE such that all analogy classes of all analogy partitions are reached with positive probability in the ABEE (see also the above discussion on the bias-variance trade-off).

5.3 Mixing ABEE with other forms of reasoning

A popular approach to interpret experimental data in early rounds of multi-player interactions is the so called level k approach (Stahl and Wilson (1995), Nagel (1995), see also Crawford and Iriberi (2007) for an application to Bayesian games). When the payoff and information structures of the opponent are known to subjects (as is often the case in lab experiments), one can iterate the best-response mapping up to some level k , assuming the level 0 behavior of the opponent corresponds to some focal strategy such as the uniform randomization over the possible actions. A level k player is best-responding to the belief that the opponent behaves according to the level $k-1$ strategy. The cognitive hierarchy model (Camerer et al. (2004)) proposes a related approach in which at level k , instead of best-responding to level $k-1$, the player best-responds to the behavior generated by a

distribution (assumed to be Poisson) over lower levels.

These models aim at capturing early rounds of interactions, and as such they do not propose equilibrium concepts in which the beliefs are related to the actual behaviors. This is a priori very different from the environments envisioned in ABEE in which beliefs are coarse but tightly related to the actual behaviors (as a result of some possibly collective underlying learning process). Despite this fundamental difference, I can think of possible ways to combine the two approaches. Specifically, I would imagine an ABEE treatment to the determination of level 0 behavior, having in mind that the focal (level 0) behavior in these models is derived from some coarse statistics about aggregate behaviors as in ABEE. From this perspective, the level 0 behaviors would be determined in an equilibrium fashion, aggregating past behaviors in the various analogy classes (as in ABEE) where the behaviors would be generated by the assumed distribution of level k agents in the population. Clearly, more work is needed to see if such a combination can bring new insights, be fruitful, and in which contexts. One limitation I see is that such an approach requires a knowledge of the payoff and information structures of the opponent (to be able to iterate the best-response), and in many field (or real-life) contexts, I would imagine such a knowledge is not directly available and has to be inferred from accessible observations (such as past behaviors).

More generally and beyond the level k formulation, one could combine the ABEE approach with the so called theory of mind as developed in philosophy and psychology in which subjects are viewed as putting themselves in the shoes of others before making their choice of strategy. For every player i , endowing other players $-i$ with beliefs formed in the ABEE sense given the analogy partition of player i would look like a natural approach to me.

In a number of situations, especially if one plays several times, one may receive feedback on one own's performance as well as feedback on others's behaviors. A question then arises as to how to combine reasoning based on feedback about others' behaviors (as in ABEE) with reasoning based on performance (as in Valuation Equilibrium or Sampling Equilibrium). Alternatively, these different modes of reasoning can be regarded as competing ones, and a player may engage into selecting among these theories, possibly using the methods described above to model the choice between competing analogy partitions.

Obviously, more work is required to make progress on these questions.

6 Conclusion

There are obviously many ways to approach the modeling of bounded rationality in games. In this survey, while I have discussed succinctly a number of these approaches, I have focused on the analogy-based expectation equilibrium. In ABEE, players are viewed as forming their expectations based on aggregate statistics describing behavioral data in their various analogy classes, and the corresponding equilibrium is viewed as a steady state of a learning process in which agents would form their expectations in this way.¹¹⁸ The ABEE framework is flexible enough to accommodate many different types of mistaken expectations, as results from the possibility to vary the analogy partitions. In my view, an appealing aspect of ABEE is the plausibility of the deliberation process attributed to the players that depends only on the aggregate statistics players are assumed to have access to. More generally and beyond ABEE, I would argue that it is desirable to impose that the deliberation process of players does not depend on aspects of the environments players are not assumed to know or have learned about. Obviously, more work can and should be done to further understand how players' analogy partitions are formed, especially when these are viewed as learning facilitating devices. This should involve theoretical, experimental and also empirical research effort. In particular, on the empirical side, I would imagine econometricians seeking to estimate the distribution of analogy partitions in the same way as they estimate the distribution of private information in the standard case assuming now that the play is governed by ABEE rather than Nash equilibrium. I hope I will have persuaded the readers that there is more to be done both in terms of broadening the scope of applications of ABEE, in terms of developing the theory, and in terms of developing empirical and experimental methods suited for ABEE.

¹¹⁸In a broad sense, viewing bounded rationality through a statistical lens agrees with the general theme pursued by Sargent (1993) in macroeconomics with an emphasis on the data pooling idea (here) rather than the dynamic programming idea (in Sargent).

References

- [1] Anderson, J. R. (1991): "The adaptive nature of human categorization," *Psychological Review* 98(3), 409-429.
- [2] Antler Y. (2019): "Multilevel Marketing: Pyramid-Shaped Schemes or Exploitative Scams?" mimeo
- [3] Antler Y. and B. Bachi (2020): "Searching Forever After" mimeo
- [4] Arad A. and A. Rubinstein (2019): "Multidimensional Reasoning in Games: Framework, Equilibrium, and Applications," *American Economic Journal: Microeconomics*, 11(3), 285-318.
- [5] Aumann R., (2008): "Rule-Rationality versus Act-Rationality," Discussion Paper Series dp497, The Federmann Center for the Study of Rationality, the Hebrew University, Jerusalem.
- [6] Axelrod R. (1984), *The Evolution of Cooperation*, Basic Books, ISBN 978-0-465-02122-2
- [7] Azariadis C. and R. Guesnerie (1982): "Prophéties créatrices et persistance des théories," *Revue économique* 33, (5), 787-806
- [8] Banerjee, A. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics* 107, 797-817.
- [9] Banerjee A., S. Merugu, I. S. Dhillon, and J. Ghosh (2005): "Clustering with Bregman Divergences," *The Journal of Machine Learning*
- [10] Barron, K., S. Huck and P. Jehiel (2020): "Everyday econometricians: Selection neglect and overoptimism when learning from others" mimeo.
- [11] Battigalli P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci (2015): "Self-Confirming Equilibrium and Model Uncertainty," *American Economic Review* 105(2): 646–677

- [12] Battigalli, Pierpaolo. 1987. "Comportamento Razionale ed Equilibrio nei Giochi e nelle Situazioni Sociali." Unpublished.
- [13] Bergemann, D. and S. Morris (2005), Robust Mechanism Design, *Econometrica* 73, 1771-1813.
- [14] Berk, R. (1966), Limiting Behavior of Posterior Distributions when the Model is Incorrect, *Annals of Mathematical Statistics* 37, 51-58.
- [15] Bianchi M. and P. Jehiel (2015): "Financial Reporting and Market Efficiency with Extrapolative Investors" *Journal of Economic Theory* **157**, 842–878.
- [16] Bianchi M. and P. Jehiel (2020): "Bundlers' dilemmas in financial markets with sampling investors" *Theoretical Economics*
- [17] Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy* 110, 992-1026.
- [18] Bilancini, E. and L. Boncinelli (2018): "Signaling to Analogical Reasoners Who Can Costly Acquire Information," *Games and Economic Behavior*, 110, 50-57.
- [19] Binmore, K. (1987): "Modeling rational players: Part 1," *Economics and Philosophy*, 3(2), 179-214.
- [20] Bohren A. (2016): "Informational Herding with Model Misspecification," *Journal of Economic Theory*, 163, 222-247.
- [21] Bohren A. and D. Hauser (2018): "Social Learning with Model Misspecification: A Framework and a Characterization," PIER Working Paper 18-017
- [22] Bordalo P., N. Gennaioli and A. Shleifer (2012): "Salience Theory of Choice Under Risk," *The Quarterly Journal of Economics*, 127 (3), 1243–1285.
- [23] Brown, G.W. (1951) "Iterative Solutions of Games by Fictitious Play" In *Activity Analysis of Production and Allocation*, T. C. Koopmans (Ed.), New York: Wiley.

- [24] Brunnermeier M. K. and J. A. Parker (2005): "Optimal Expectations," *American Economic Review*, 95(4), 1092-1118.
- [25] Camerer C. and T. H. Ho (1999): "Experience-weighted Attraction Learning in Normal Form Games," *Econometrica* 67(4), 827-874.
- [26] Camerer CF, TH Ho and JK Chong (2004): "A cognitive hierarchy model of games," *Quarterly Journal of Economics* 119 (3), 861-898.
- [27] Chen H-C, J-F Thisse and J. Friedman (1997): "Boundedly Rational Nash Equilibrium: A Probabilistic Choice Approach," *Games and Economic Behavior* 18, 32-54.
- [28] Compte O. and A. Postlewaite (2018): *Ignorance and Uncertainty*, Econometric Society Monograph
- [29] Conlinsk J. (1996): "Why Bounded Rationality?," *Journal of Economic Literature*, 34(2), 669-700.
- [30] Crawford V. P. and N. Iriberry (2007): "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?," *Econometrica* 75 (6), 1721-1770
- [31] Crawford V., and J. Sobel (1982): 'Strategic information transmission,' *Econometrica* 50, 1431–1451.
- [32] Dal Bó E., P. Dal Bó and E. Eyster (2018): "The Demand for Bad Policy when Voters Underappreciate Equilibrium Effects," *The Review of Economic Studies*, 85(2), 964–998.
- [33] Danz D., S. Huck and P. Jehiel (2016): "Public statistics and private experience: Varying feedback information in a take-or-pass game" *German Economic Review*, 17 (3), 359-377 (special issue in honor of Reinhard Selten).
- [34] Dekel E., J. C. Ely and O. Yilankaya (2007): "Evolution of Preferences," *The Review of Economic Studies*, 74 (3), 685–704.

- [35] Dekel, E., Fudenberg, D., Levine, D.K. (1999): "Payoff information and self-confirming equilibrium," *Journal of Economic Theory* 89, 165-185.
- [36] Dekel, E., Fudenberg, D., Levine, D.K. (2004): "Learning to play Bayesian games," *Games Econ. Behav.* 46, 282–303.
- [37] Eliaz, K. and R. Spiegler (2020), "A Model of Competing Narratives," forthcoming *American Economic Review*.
- [38] Eliaz K., R. Spiegler, and H. C. Thyssen (2020): "Strategic interpretations," mimeo.
- [39] Ellis A. and M. Piccione (2017): "Correlation misperception in choice," *American Economic Review* 107 (4), 1264-92
- [40] Embrey, M., G. R. Fréchet, and S. Yuksel (2018): "Cooperation in the Finitely Repeated Prisoner's Dilemma," *Quarterly Journal of Economics*, 133(1): 509-551.
- [41] Esponda, I., (2008): "Behavioral equilibrium in economies with adverse selection," *Am. Econ. Rev.* 98, 1269–1291.
- [42] Esponda I. and D. Pouzo (2016): "Berk–Nash equilibrium: A framework for modeling agents with misspecified models," *Econometrica*, 84(3):1093–1130.
- [43] Esponda, I., and E. Vespa. (2018): "Endogenous Sample Selection: A Laboratory Study," *Quantitative Economics* 9(1), 183-216.
- [44] Ettinger, D and P. Jehiel (2005): "Towards a Theory of Deception," PSE working paper.
- [45] Ettinger, David, and Philippe Jehiel. 2010. "A Theory of Deception." *American Economic Journal: Microeconomics*, 2 (1): 1-20.
- [46] Ettinger D. and P. Jehiel (2020): 'An experiment on deception, reputation, and trust,' forth. *Experimental Economics*.
- [47] Eyster, E. and M. Piccione (2013): "An approach to asset pricing under incomplete and diverse perceptions," *Econometrica* 81(4), 1483-1506.

- [48] Eyster, E., and M. Rabin (2005): "Cursed Equilibrium," *Econometrica* 73 (5): 1623–72.
- [49] Eyster, E., Rabin, M. (2010): "Naïve Herding in Rich-Information Settings," *American Economic Journal: Microeconomics*, 2(4), 221-243.
- [50] Friedman, E. (2019): "Stochastic Equilibria: Noise in Actions or Beliefs?" mimeo
- [51] Friedman J. and C. Mezzetti (2005): "Random belief equilibrium in normal form games," *Games and Economic Behavior*, 51 (2), 296-323.
- [52] Fryer R, and M. Jackson (2008): "A Categorical Model of Cognition and Biased Decision-Making," *Contributions in Theoretical Economics*. 2008;8 (1).
- [53] Fudenberg, D and Y. Kamada (2018): "Rationalizable partition-confirmed equilibrium with heterogeneous beliefs," *Games and Economic Behavior* 109, 364-381.
- [54] Fudenberg D. and D. M. Kreps (1995): "Learning in Extensive Games," I: Self-Confirming and Nash Equilibrium. *Games and Economic Behavior*, 8, 20-55.
- [55] Fudenberg D., G. Lanzani and P. Strack (2020): "Limits Points of Endogenous Misspecified Learning" mimeo
- [56] Fudenberg D. and D. K. Levine (1993): "Self-Confirming Equilibrium," *Econometrica*, 61:523-545.
- [57] Fudenberg D. and D. Levine. *Learning in Games*. M.I.T. Press, 1998.
- [58] Fudenberg D. and D. K. Levine (1989): "Reputation and equilibrium selection in games with a patient player," *Econometrica*, 57(4): 759–778.
- [59] Fudenberg D. and A Peysakhovich (2016): "Recency, records, and recaps: Learning and nonequilibrium behavior in a simple decision problem," *ACM Transactions on Economics and Computation (TEAC)* 4 (4), 1-18
- [60] Gabaix X. (2014): "A Sparsity-Based Model of Bounded Rationality," *Quarterly Journal of Economics*, 129(4), 1661–1710.

- [61] Gagnon-Bartsch T., M. Rabin, and J. Schwartzstein (2020): "Channeled attention and stable errors," mimeo.
- [62] Garcia-Pola B., J. Kovarik and I. Nagore (2020): "Non-equilibrium Play in Centipede Games," *Games and Economic Behavior*, 120, 391–433.
- [63] Gärdenfors, P. (2000): *Conceptual Spaces: The Geometry of Thought*, MIT Press, Cambridge, MA.
- [64] Gigerenzer, G. (2001): "The adaptive toolbox," In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (p. 37–50). The MIT Press.
- [65] Gilboa I. and D. Schmeidler (1995): "Case-based decision theory," *Quart. J. Econ.* 110, 605–639.
- [66] A. Guarino and P. Jehiel (2013): "Social learning with coarse inference," *American Economic Journal: Microeconomics* 5(1), 147-74.
- [67] Grimée, J. (2018): "Bounded Rationality in Local Tax Competition," Master thesis APE.
- [68] Grimm V. and F. Mengel (2012): "An experiment on learning in a multiple games environment," *Journal of Economic Theory*, 147(6), 2220-2259.
- [69] Hagenbach J. and F. Koessler (2019): "Cheap talk with coarse understanding," forthcoming *Games and Economic Behavior*.
- [70] Hagenbach J. and F. Koessler (2017): "The Streisand effect: Signaling and partial sophistication," *Journal of Economic Behavior & Organization*, 143: 1-8.
- [71] Haile P. A. , A. Hortaçsu, and G. Kosenok (2008): "On the Empirical Content of Quantal Response Equilibrium," *American Economic Review*, 98, 180–200
- [72] Harrison, J. M. and D. M. Kreps (1978): "Speculative investor behavior in a stock market with heterogeneous expectations," *Quarterly Journal of Economics* 92(2), 323-36.

- [73] Harsanyi, J.C (1967-1968): "Games With Incomplete Information Played by Bayesian Players. Parts I, II, III," *Management Science* 14 (3, 5, 7).
- [74] Hart S. (2005): "Adaptive heuristics," *Econometrica* 73 (5), 1401-1430.
- [75] Hart S. and A. Mas-Colell (2003): "Uncoupled dynamics do not lead to Nash equilibrium," *American Economic Review* 93 (5), 1830-1836.
- [76] Heifetz A. and D. Samet (1988): "Knowledge spaces and arbitrarily high rank," *Games and Economic Behavior* 22, 260-273.
- [77] Heckman, J. (1979): "Sample Selection Bias as a Specification Error," *Econometrica* 47 (1): 153–61.
- [78] Heller, Y. and E. Winter (2020): "Biased-Belief Equilibrium," *American Economic Journal: Microeconomics*, 12(2), 1-40.
- [79] Holland, J. (1975): *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence*. Cambridge, MA: MIT Press.
- [80] Huck S., P. Jehiel and T. Rutter (2011): "Learning Spillover and Analogy-based Expectations: A Multi-Game Experiment," *Games and Economic Behavior* 71(2) (2011), 351-365.
- [81] Jackson M. and E. Kalai (1997): "Social learning in recurring games," *Games and Economic Behavior* 21 (1–2), 102-134.
- [82] Jaynes, E. T. (1957). "Information Theory and Statistical Mechanics" (PDF). *Physical Review. Series II*. 106 (4): 620–630. Bibcode:1957PhRv..106..620J.
- [83] Jehiel P. (1995): 'Limited horizon forecast in repeated alternate games,' *Journal of Economic Theory* 67, 497-519.
- [84] Jehiel P. (1998): 'Learning to play limited forecast equilibria,' *Games and Economic Behavior* 22, 274-298.

- [85] Jehiel P. (1999): "Predicting by Analogy and Limited Foresight in Games" CiteSeer
- [86] Jehiel P. (2001): 'Limited foresight may force cooperation' *Review of Economic Studies* **68**, 369-391.
- [87] Jehiel P. (2005): 'Analogy-based expectation equilibrium,' *Journal of Economic Theory* **123**, 81-104.
- [88] Jehiel, P. (2011) "Manipulative Auction Design" *Theoretical Economics* **6**(2), 185-217.
- [89] Jehiel P. (2015): "On transparency in organizations," *Review of Economic Studies* **82** (2), 736-761
- [90] Jehiel P. (2018): "Investment strategy and selection bias: An equilibrium perspective on overoptimism," *American Economic Review* **108**(6), 1582–1597.
- [91] Jehiel P. (2019): "Robust Mechanism Design: An Analogy-Based Expectation Equilibrium Perspective" in: Laslier J.F., Moulin H., Sanver M., Zwicker W. (eds) *The Future of Economic Design. Studies in Economic Design*. Springer, Cham
- [92] Jehiel P. (2021): 'Communication with forgetful liars', *Theoretical Economics*
- [93] Jehiel P. and F. Koessler (2008): 'Revisiting Bayesian games with analogy-based expectations,' *Games and Economic Behavior* **62**, 533-557.
- [94] Jehiel P. and L. Lamy (2015) "On absolute auctions and secret reserve prices" *Rand Journal of Economics* **46** (2) (2015), 241–270.
- [95] Jehiel, P., M. Meyer-ter-Vehn, B. Moldovanu and W.R. Zame (2006), The Limits of Ex-Post Implementation, *Econometrica* **74**, 585-610
- [96] Jehiel P. and K. Mierendorff (2021): "Auction Design with Data-Driven Misspecifications" mimeo
- [97] Jehiel P and E. Mohlin (2021): "Cycling and Categorical Learning in Decentralized Adverse Selection Economies," mimeo.

- [98] Jehiel, P. and D. Samet (2005): "Learning to play games in extensive form by valuation," *Journal of Economic Theory*, 124, 129–148
- [99] Jehiel P. and D Samet (2007): "Valuation Equilibrium," *Theoretical Economics* **2**, 163–185.
- [100] Jehiel P. and L. Samuelson (2012): "Reputation with analogical reasoning," *Quarterly Journal of Economics* 127, 1927-1969
- [101] Jehiel P. and J. Singh (2021): 'Multi-state choices with aggregate feedback on unfamiliar alternatives' *Games and Economic Behavior*.
- [102] Jehiel P. and G. Weber (2021): "Clustering and Analogy-Based Expectation Equilibrium with an Application to Product-Differentiated Oligopoly," mimeo
- [103] Liang, A. (2020): "Games of Incomplete Information Played By Statisticians," mimeo.
- [104] Kalai E. (2004): "Large Robust Games," *Econometrica*, 72(6), 1631-1665
- [105] Kalai E. and E. Lehrer (1993): "Rational Learning Leads to Nash Equilibrium," *Econometrica* 61 (5), . 1019-1045
- [106] Ke S. (2019): "Boundedly rational backward induction," *Theoretical Economics*,14(1), 103-134
- [107] Klibanoff P., M. Marinacci, and S. Mukerji (2005): "A smooth model of decision making under ambiguity," *Econometrica* 73 (6), 1849-1892
- [108] Koehler, Jonathan J., and Molly Mercer (2009): "Selection Neglect in Mutual Fund Advertisements." *Management Science* 55 (7): 1107–21.
- [109] Köszegi B. and A. Szeidl (2013): "A model of focusing in economic choice," *The Quarterly Journal of Economics* 128 (1), 53-104
- [110] D. Kreps and R. Wilson (1982): "Sequential equilibria," *Econometrica* 50, 863–894.

- [111] Krockow E. M., , B. D. Pulford and A. M. Colman (2018): "Far but finite horizons promote cooperation in the Centipede game ," *Journal of Economic Psychology* 67, 191-199.
- [112] Kurz, M. (1994): "On rational belief equilibria," *Economic Theory* 4, 859-876.
- [113] Lasry, J.-M. and P-L Lions (2006): "Jeux à champ moyen. II – Horizon fini et contrôle optimal" [Mean field games. II – Finite horizon and optimal control]. *Comptes Rendus Mathématique*, 343 (10): 679–684.
- [114] Lee, I. H., (1993): "On the Convergence of Informational Cascades," *Journal of Economic Theory*, 61, 395-411.
- [115] Lehrer, E. (2012): "Partially Specified Probabilities: Decisions and Games," *American Economic Journal: Microeconomics* 4 (1): 70–100.
- [116] Lettau M. and H. Uhlig (1999): "Rules of Thumb versus Dynamic Programming," *American Economic Review*. 89(1), 148-174.
- [117] Levy G. and R. Razin (2019): "Echo Chambers and Their Effects on Economic and Political Outcomes," *Annual Review of Economics* Volume 11, 303-328.
- [118] Li Calzi, M (1995): "Fictitious Play by Cases," *Games and Economic Behavior*, 11, 64-89.
- [119] Lipman B. L. (1991): "How to Decide How to Decide How to . . . : Modeling Limited Rationality," *Econometrica*, 59(4), 1105-25.
- [120] Madarasz K. (2012): "Information Projection: Model and Applications," *The Review of Economic Studies*, 79(3), 961-985
- [121] Mailath G. and L. Samuelson (2020): "Learning under Diverse World Views: Model Based Inference" *American Economic Review*, forthcoming
- [122] Marimon, R.; McGrattan, E. and Sargent, T. J. "Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents." *Journal of Economic Dynamics and Control*, May 1990, 14(2), 329–73.

- [123] McAfee R.P. (1993): "Mechanism design by competing sellers," *Econometrica*, 61(6):1281–1312.
- [124] McFadden, D. (1973): "Conditional Logit Analysis of Qualitative Choice Behavior," in P. Zarembka, ed., *Frontiers in Econometrics* (New York: Academic Press), 105–42.
- [125] R. Mc Kelvey and T. Palfrey (1995): "Quantal response equilibrium for normal form games," *Games Econ. Behav.* 10, 6–38.
- [126] Mengel, F. (2012): "Learning Across Games," *Games and Economic Behavior*, 74(2), 601-619.
- [127] Mengel, F. (2014): "Learning by (limited) forward looking players," *Journal of Economic Behavior and Organization* 108, 59-77.
- [128] Mertens, J.-F. and S. Zamir (1985): "Formulation of Bayesian analysis for games with incomplete information," *Int J Game Theory* 14, 1–29.
- [129] Miyasawa, K. (1961): "On the convergence of the learning process in a 2x2 non-zero-sum game," *Economic Research Program, Princeton University, Research Memo* 33.
- [130] Miettinen T. (2009): "The partially cursed and the analogy-based expectation equilibrium," *Economics Letters*, 105(2), 162-164.
- [131] Milgrom, P. (1981): 'Good News and Bad News: Representation Theorems and Applications,' *Bell Journal of Economics* 12 (2): 380–91.
- [132] Milgrom, P., Stockey, N., 1982. Information, trade and common knowledge. *J. Econ. Theory* 26, 17–27.
- [133] Mohlin, E. (2014) "Optimal Categorization" *Journal of Economic Theory* 152, pp. 356–381
- [134] Moinas S. and S. Pouget (2013): "The Bubble Game: An Experimental Study of Speculation," *Econometrica* Volume 81, 1507-1539.

- [135] Mongin P. and B. Walliser (1988): "Infinite Regressions in the Optimizing Theory of Decision," in *Risk, decision and rationality*. Ed.: B. R. Munier. Dordrecht: D. Reidel, pp. 435-57.
- [136] Mullainathan, S. Schwartzstein, J. and A. Shleifer (2008): "Coarse Thinking and Persuasion," *Quarterly Journal of Economics*, 123, 577-619.
- [137] Nagel, R. (1995): 'Unraveling in Guessing Games: An Experimental Study', *American Economic Review*, 85, 1313-1326
- [138] Neyman A (1985): "Bounded complexity justifies cooperation in the finitely repeated prisoners' dilemma," *Economics letters* 19 (3), 227-229
- [139] Olea J. L., P. Ortoleva, M. M Pai, and A. Prat (2019): "Competing models," mimeo.
- [140] M. Osborne, A. Rubinstein (1998): "Games with procedurally rational players," *Amer. Econ. Rev.* 88, 834–847.
- [141] Osborne, M. J. and Rubinstein, A. (2003): "Sampling equilibrium, with an application to strategic voting," *Games and Economic Behavior*, 45(2), 434–441.
- [142] Pearl, J. (2009), *Causality: Models, Reasoning and Inference*, Cambridge University Press, Cambridge.
- [143] Piccione M. and A. Rubinstein (1997): "On the interpretation of decision problems with imperfect recall," *Games Econ. Behav.* 20, 3–24.
- [144] 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, 212-233.
- [145] R. Radner (1980): "Collusive behavior in non-cooperative epsilon equilibria of oligopolies with long but finite lives," *Journal of Economic Theory*, 22, 121–157.
- [146] Rampal J. (2017): "Limited foresight equilibrium," Technical report.
- [147] R. Rosenthal (1982): "Games of perfect information, predatory pricing and the chain-store paradox," *J. Econ. Theory* 25, 83–96.

- [148] R. Rosenthal (1993): "Rules of Thumb in Games," *Journal of Economic Behavior and Organization*, 22, 1-13.
- [149] Ross, L. (1977): "The Intuitive Psychologists and his Shortcomings," in Berkowitz L. (ed), *Advances in Experimental Social Psychology* (Vol 10, 173-220), New York: Academic.
- [150] Rubinstein A. (1986): "Finite automata play the repeated prisoner's dilemma," *Journal of Economic Theory*, 39 (1), 83-96.
- [151] Rubinstein, A. (1988): "Similarity and decision-making under risk (is there a utility theory resolution to the Allais paradox?)," *J. Econ. Theory* 46, 145–153.
- [152] Rubinstein, A. (1989): "The electronic mail game: Strategic behavior under 'almost common knowledge' *Amer. Econ. Rev.* 79, 385–391.
- [153] Rubinstein, A. (1998): *Modeling Bounded Rationality*, MIT Press.
- [154] Rubinstein A. and A. Wolinsky (1994): "Rationalizable Conjectural Equilibrium: Between Nash and Rationalizability," *Games and Economic Behavior* 6, 299-311.
- [155] Salant Y. and J. Cherry (2020): "Statistical Inference in Games ," *Econometrica* forthcoming.
- [156] Samuel A.L. (1959): "Some studies in machine learning using the game of checkers," *IBM J. Res. Devel.* 3, 210–229.
- [157] Samuelson L. (2001): "Analogies, adaptation, and anomalies," *J. Econ. Theory* 97, 320–366.
- [158] Sargent, T. J. (1993): *Bounded rationality in macroeconomics*. Oxford [England]: Clarendon Press.
- [159] Savage, L. J. (1954): *The Foundation of Statistics*, John Wiley & Sons, Inc.
- [160] Schank, Roger (1982): *Dynamic Memory: A Theory of Learning in Computers and People*. New York: Cambridge University Press.

- [161] Scheinkman, J. A. and Xiong, W. (2003): "Overconfidence and speculative bubbles," *Journal of Political Economy* 111(6), 1183-1219.
- [162] Schumacher, H. and H. Thysen (2018), *Equilibrium Contracts and Boundedly Rational Expectations*, mimeo.
- [163] Schwartzstein J. (2014): "Selective Attention and Learning," *Journal of European Economic Association* 12 (6), 1423-1452
- [164] Sebenius, J.K., and J. Geanakoplos (1983): "Don't bet on it: Contingent agreements with asymmetric information," *J. Amer. Statist. Assoc.* 78, 424–426.
- [165] Selten, R. (1975). "A Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games". *International Journal of Game Theory.* 4 (1): 25–55
- [166] R. Selten, R. Stoecker (1986): "End behavior in sequences of finite prisoner's dilemma supergames," *J. Econ. Behav. Organ.* 7, 47–70.
- [167] Simon, Herbert A. (1956): "Rational Choice and the Structure of the Environment" *Psychological Review.* 63 (2): 129–138.
- [168] Sims C. (2003): "Implications of rational inattention," *Journal of Monetary Economics* 50 (3), 665-690
- [169] Smith, L., and P. Sørensen (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371-398.
- [170] Smith V. L. , G. L. Suchanek and A. W. Williams (1988): "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica* 56(5), 1119-1151
- [171] Sobel, J. (1985): 'A theory of Credibility', *Review of Economic Studies*, 52, 557-573.
- [172] Spiegel, R. (2006a): "Competition over agents with boundedly rational expectations," *Theoretical Economics*, 1(2), 207–231.
- [173] Spiegel, R. (2006b): "The market for quacks. *Review of Economic Studies*," 73(4), 1113–1131.

- [174] Spiegel, R. (2016): "Bayesian Networks and Boundedly Rational Expectations," *Quarterly Journal of Economics* 131, 1243-1290.
- [175] Spiegel, R. (2017): "Data Monkeys: A Procedural Model of Extrapolation From Partial Statistics," *Review of Economic Studies* 84, 1818-1841.
- [176] Spiegel, R. (2019): "Behavioral Economics and the Atheoretical Style," *American Economic Journal: Microeconomics*, 11, 173-94.
- [177] Spiegel R. (2020): "Behavioral Implications of Causal Misperceptions," *Annual Review of Economics*, forthcoming.
- [178] Stahl, D. and Wilson, P. (1995): 'On Player's Modals of other Players: Theory and Experimental Evidence', *Games and Economic Behavior*, 10, 218-254.
- [179] Steiner J. and C. Stewart (2008): "Contagion through learning," *Theoretical Economics* 3 (4), 431-45
- [180] Tirole, J. (1982): "On the Possibility of Speculation Under Rational Expectations," *Econometrica*, 50 (5), 1163-1182.
- [181] Watkins C. and P. Dayan (1992): "Q-learning," *Machine Learning*, 8, 279-292.
- [182] Weibull JW (1997): *Evolutionary game theory*, MIT press
- [183] Young HP (2004): *Strategic learning and its limits*, Arne Ryde Memorial Lectures.