

Modeling Bounded Rationality in Economic Theory: Four Examples *

Ariel Rubinstein

School of Economics, Tel Aviv University
and Department of Economics, New York University

**Published in Handbook of Bounded Rationality, edited by
Riccardo Viale, Routledge, (2021), 423-435.**

Revised: August, 2026.

ABSTRACT: The paper provides an introduction to the growing literature in Economics on models of Bounded Rationality, in which reasoning procedures are the cornerstones of the model. Four models are presented, each of which demonstrates a different type of Bounded Rationality:

(i) Limited ability to solve a set of propositions: A mechanism design problem is investigated in an environment where agents try to cheat effectively using a procedure which is anchored in the truth.

(ii) Belief formation on the basis of a small sample: A model of manipulative voting is presented in which a voter forms beliefs about his vote being pivotal based on a small sample of observations of other voters' intentions.

(iii) Reducing the complexity of strategies: Each player in a repeated Prisoner's Dilemma desires not only to increase his payoff but also to reduce the complexity of his strategy.

(iv) Diverse views of the world: A seller exploits differences in buyers' abilities to recognize time series patterns in order to increase his profits beyond what they would be with a fixed price.

*I would like to thank Jacob Glazer, Martin Osborne, Michele Piccione and Rani Spiegler for their comments.

1. Introduction

The terms *Bounded Rationality* and *Economic Theory* mean different things to different people. For me (see Rubinstein (2012)), Economic Theory is a collection of stories, usually expressed in formal language, about human interactions that involve joint and conflicting interests. Economic Theory is not meant to provide predictions of the future. At most, it can clarify concepts and provide non-exclusive explanations of economic phenomena. In many respects, a model in Economic Theory is no different than a story. Both a story and a model are linked to reality in an associative manner. Both the storyteller and the economic theorist have in mind a real-life situation but do not consider the story or the model to be a full description of reality. Both leave it to the reader to draw their own conclusions, if any.

Models of bounded rationality are for me (see Rubinstein (1998)) models that include explicit references to procedural aspects of decision-making, which are crucial for the derivation of the analytical results. A common critique of bounded rationality models is that they are more specific, less general, and more arbitrary than models in the mainstream areas of Economic Theory, such as general equilibrium or game theory, in which full rationality is assumed. In response, I would claim that every model makes very (very) special assumptions. Without strong assumptions, there would be no conclusions. It is true that rationality is a special assumption, since it is viewed by many as being normative, whereas models of bounded rationality are viewed as dealing with deviations from normative behavior. Returning to the analogy of a story: What story is more interesting - one about people who behave “according to the book” or one about people who deviate from normative behavior?

Not every model that is inconsistent with some aspect of rationality is a model of bounded rationality. A model in which agents violate a rationality requirement, without specifying a reasoning procedure that accounts for the violation, is a bad model rather than a model of bounded rationality. A good model of bounded rationality should include a procedure of reasoning that makes sense and is somewhat related to what we observe in real life.

I am not a big fan of abstract methodological discussions. I prefer to demonstrate an approach by discussing examples. Accordingly, the article discusses four models in

which economic agents are assumed to reason in systematic ways using a well-defined procedure that is outside the standard scope of rational behavior. My choice of models is totally subjective. I am partial to these models because I was involved in constructing and analysing them over the last four decades. Rubinstein (1998) surveys the field at the time, while Spiegler (2011) surveys more recent models in which aspects of bounded rationality have been inserted into classic models of Industrial Economics.

2. Bounded Rationality and Mechanism Design

Story: The director of a prestigious MBA program has been persuaded by Choi, Kariv, Müller, and Silverman (2014) that transitivity is strongly correlated with success in life (as measured by wealth). Thus, he decides to accept only those candidates who hold transitive preferences.

Accordingly, the director designs a simple test. He presents three alternatives a, b , and c to each candidate and asks them to respond to a questionnaire consisting of three questions $Q(a, b)$, $Q(b, c)$, and $Q(a, c)$, where $Q(x, y)$ is the quiz question:

Do you prefer x to y or y to x ?

☐ I prefer x to y .

☐ I prefer y to x .

For each of the three questions, a candidate must respond by clicking on one and only one of the two possible answers. Thus, the questionnaire has eight possible sets of responses.

The director is obligated by law to inform candidates about the conditions under which they will be admitted to the program. The director has specified the following two conditions:

R1: If you prefer a to b and b to c then you must prefer a to c .
R2: If you prefer c to b and b to a then you must prefer c to a .

The candidates are reminded that an “if” proposition is violated only if its antecedent (the “if” part) is satisfied and its consequent (the “then” part) is not.

The director hopes that the candidates will feel obliged to report the truth and that the simple questionnaire will therefore perfectly separate good candidates, who hold

transitive preferences, from bad candidates, who do not. In order to encourage instinctive responses, the director also sets a short time limit for completing the questionnaire.

A disappointment: the director learns that all the candidates have been admitted and concludes that the candidates with cyclical preferences have gamed the system. The frustrated director opens an investigation. Researchers are rushed to the scene. They interview candidates to find out how they had managed to answer the questionnaire successfully. Apparently, candidates treated the questionnaire as a puzzle to be solved in order to be admitted to the program. The time limit made solving the puzzle a challenging task. The following procedure was identified:

Step 1: Examine whether your honest set of answers satisfies all the conditions.

If it does, then happily submit those answers. If not, go to Step 2.

Step 2: Find a condition that is violated by your honest answers (that is, your true answers satisfy the antecedent but not the consequent of the condition). Try changing only the answer that appears in the consequent. If the modified set of answers satisfies all conditions, then submit them. If not, repeat Step 2, beginning each time with your honest answers, until all the violated conditions have been considered. Then, proceed to Step 3.

Step 3: Give up. Submit the honest answers (and be rejected).

Faced with the questionnaire and the admission conditions, a good candidate (who has transitive preferences) reports the truth and is admitted. What about a bad candidate with cyclical preferences $a \succ b \succ c \succ a$? An honest response satisfies R2 (the antecedent is false) but violates R1. The candidate retains his truthful answers to $Q(a, b)$ and $Q(b, c)$ and modifies his answer to $Q(a, c)$. Thus, he responds as if he holds the preferences $a \succ b \succ c$ and is admitted. Similarly, a candidate with cyclical preferences $c \succ b \succ a \succ c$ also changes the answer to $Q(a, c)$, but responds as if he holds the preferences $c \succ b \succ a$. In this way, all the candidates are admitted.

Realizing that people are prepared to cheat (when they have to) by using the procedure to find a persuasive set of responses, the director tries to come up with a modified set of admission conditions so that all the good candidates will pass the test while the bad candidates will not. He adds two new admission conditions, so that the new set of admission conditions is as follows:

- R1: If you prefer a to b and b to c , then you must prefer a to c .
R2: If you prefer c to b and b to a , then you must prefer c to a .
R3: If you prefer a to b and a to c , then you must prefer c to b .
R4: If you prefer c to a and c to b , then you must prefer a to b .

This set of conditions seems a bit odd. How would the candidates respond to it?

First consider a good candidate who holds one of the transitive relations $b \succ a \succ c$, $b \succ c \succ a$, $a \succ c \succ b$, or $c \succ a \succ b$. None of the propositions is violated, and responding truthfully leads to acceptance, as desired by both the candidate and the director.

What about a good candidate who holds the preferences $a \succ b \succ c$ (or $c \succ b \succ a$)? If he responds truthfully, then he violates R3 (or R4) and is rejected. However, he is induced by the violation of R3 (or R4) to modify his answer to $Q(b, c)$ (or $Q(a, b)$); this leads him to respond as if he holds the preferences $a \succ c \succ b$ (or $c \succ a \succ b$) and thus he cheats successfully.

On the other hand, a bad candidate with cyclical preferences $a \succ b \succ c \succ a$ (or a candidate with the reverse cycle $c \succ b \succ a \succ c$) violates only R1 (or R2) and is led to try only the set of answers corresponding to the preferences $a \succ b \succ c$ (or $c \succ b \succ a$) which does not satisfy the set of conditions. Thus, such a candidate fails the test, as desired by the program director.

A happy ending, at least for the program director. Interestingly, he does not mind that some candidates cheat. The questionnaire “works” in the sense that it separates the good candidates from the bad candidates. Actually, the director also finds that there is no alternative set of conditions that will separate the two types of candidates without some good candidates having to cheat. He concludes that white lies are sometimes necessary.

Discussion: This section is based on Glazer and Rubinstein (2012). The model comes under the rubric of bounded rationality since it explicitly specifies a process of reasoning used by the individuals. The paper characterizes the circumstances under which the designer of the mechanism is able to implement his target given that the individuals use the procedure discussed above. Conditions under which implementation does not require any good candidate to cheat are established as well.

No one claims that most people use this exact procedure. Nonetheless, Glazer and Rubinstein (2012) provide experimental evidence that the two central ingredients of the procedure are present in many people's minds: (i) The truth is the anchor for cheating. When one invents a false set of answers, he starts from the truth and modifies the answers to make them look better. (ii) When an "if" sentence is violated, an individual who wishes to cheat successfully will modify his true set of answers by reversing the relevant answer so that the consequent becomes true. Note that without (i), separating the good candidates from the bad candidates would be impossible.

In the current story, the bounded rationality of the candidates is an advantage for the designer. If all candidates were fully rational, as is commonly assumed in the mechanism design literature, all of them would be able to game the system if necessary. It is the cognitive imperfection of the individuals that *opens the door* to obtaining a desired outcome.

Bibliographic comments:

(i) Glazer and Rubinstein (2014) study a related model in which a candidate responds to a questionnaire without being notified of the acceptance conditions, though he has access to the data about the set of responses that lead to admission and tries to make sense of things. There is a bound on the complexity of the regularities that the candidates can detect in the data. It is shown that whatever this bound is, the director can construct a sufficiently complex questionnaire so that agents who respond honestly to the questionnaire will be treated optimally and the probability that a dishonest agent will cheat successfully is "very small".

(ii) de Clippel (2014) investigates Maskin's classical implementation question in environments in which individuals follow systematic procedures of choice that are not consistent with rational behavior.

(iii) Li (2017) is motivated by an observation that bidders in an ascending-bid auction tend to wait until the bid reaches their reservation value, whereas a majority of bidders do not report their true values in the "equivalent" second-price auction. An *obviously dominant strategy* is one in which the decision at each decision node is "obvious" in the sense that it does not depend on conjecture about the other players' actions. More precisely, at every relevant decision point, the worst possible outcome from following the prescribed strategy is at least as good as the best possible outcome from

deviating. Li identifies several targets that are implementable using an extensive-form game in which the desired outcome is obtained through obviously dominant strategies. Glazer and Rubinstein (1996) is an earlier paper that suggests a different criterion for simple implementation.

3. Elections and Sampling Equilibrium

Story: Three candidates are running for office. A large number of voters will participate in the election and each will cast one vote in favor of one of the candidates. The issues on the agenda are hot and all of the voters are expected to cast their ballots. The winner will be the candidate with the largest number of votes (even if he does not achieve an absolute majority).

The candidates are labeled as Left, Center and Right. The population of voters is split into three classes *Leftist*, *Centrist* and *Rightist* with the proportions λ , μ and ρ :

Leftist: Individuals with the preferences $L \succ C \succ R$.

Centrist: Individuals with the preferences $C \succ L \sim R$.

Rightist: Individuals with the preferences $R \succ C \succ L$.

It is known that the Leftist and Rightist groups are roughly of the same size and that the Centrist group is smaller but not negligible ($1/3 < \lambda = \rho < 3/7$). Some observers predict that C will win since some of the leftists and rightists will vote for C in order to prevent the election of the candidate at the other extreme. Some view such a result as a desirable compromise. Others find it unacceptable that the least popular candidate might win the election.

Voters by nature tend to vote sincerely, but they will nonetheless vote for another candidate if they think that their candidate will certainly lose and the race between the other two candidates is close. In such an event, they vote for their second-best candidate.

Surveys are forbidden and voters base their prediction of the election results on occasional discussions with casual acquaintances who express their voting intentions (without specifying their first choice). As is often the case, people say that they “talked with several people” but actually mean that they talked with just two. Thus, based on

a sample of size two, people decide who to vote for. Therefore, they may change their minds during the campaign according to the sample results.

Some game theorists and political scientists analyse complicated models in which voters think strategically, in the sense that they put themselves in the shoes of other voters, calculate correctly the distribution over states conditional on the infinitesimal probability event that their personal vote will be pivotal and vote optimally given that event. No traces of such behavior can be found in our story (nor in real life). No standard game-theoretical considerations are used. At each moment in the campaign, each voter constructs his prediction and decides who to vote for based on two independent random draws from the current population distribution.

Accordingly, if a voter samples two people who intend to vote for the same candidate, the voter treats the election as already decided, concludes that his vote will not make any difference and votes sincerely. If the two sampled individuals intend to vote for two different candidates, the voter views himself as pivotal and votes for the candidate he prefers of these two (who is not necessarily his favorite). Thus, a C supporter intends to vote sincerely regardless of the sample findings. An L supporter will vote sincerely unless he draws a sample of C and R and then votes for C . An R supporter will vote for C if his sample is C and L and sincerely otherwise.

Thus, although voters are stubborn in their political views and never change their basic position, they may change their intention during the campaign on the basis of the sample. We look for a distribution of voting intentions that is stable under this sampling procedure; we call such a distribution an equilibrium.

Given that all C supporters vote C , L supporters will never vote R and R supporters will never vote L , a distribution of votes is characterized by two numbers: a , the proportion of the L camp that votes C , and b , the proportion of the R camp that votes C . In equilibrium, those proportions are stable. Notice that even when equilibrium is reached, members of the L and R camps may still change their intentions; however, in equilibrium, the changes are balanced and the distribution of votes remains constant.

Election day arrives. The polls close. Will C win? Let's examine the equilibrium. Recall that the distribution of the L, C, R groups is $(\lambda, 1 - 2\lambda, \lambda)$. If at some point in time the proportions of L - and R -supporters who switch to C are a and b , then the proportions of voters' intentions (following the above procedure) are $(\lambda(1 - a), \lambda a +$

$(1 - 2\lambda) + \lambda b, \lambda(1 - b)$). The proportion of L supporters who will sample C and R is $2[\lambda a + (1 - 2\lambda) + \lambda b][\lambda(1 - b)]$. The distribution is stable if this proportion is equal to a (and analogously for b). Thus, equilibrium is characterized by two equations:

$$a = 2[\lambda a + (1 - 2\lambda) + \lambda b][\lambda(1 - b)] \text{ and } b = 2[\lambda a + (1 - 2\lambda) + \lambda b][\lambda(1 - a)].$$

This system of two equations has a unique solution which can easily be calculated and accordingly, given the constraints on the parameters, the minority candidate C will win! Incidentally, if the size of C 's base dropped below $1/7$, candidate C would lose, at least according to the story we have told.

Discussion: This section is based on Osborne and Rubinstein (2003). The bounded rationality component of the model is the procedure used by each voter to estimate the results. As mentioned before, the procedure of choice assumed here is in sharp contrast to what is assumed in standard game-theoretical voting models, in which a voter is assumed to fully calculate the conditions under which his vote will be pivotal given the correct expectations about the votes of all other voters.

Bibliographic comments:

(i) The model was inspired by Osborne and Rubinstein (1998), who proposed the concept of S-1 equilibrium and applied it to symmetric two-player games. A population of players is randomly matched in pairs to play a game. A newcomer to the population samples each of the potential actions once by interviewing one person whose experience depends on the behavior of the randomly matched player he played against. Then, the newcomer chooses the action which, according to the sample, yields the best outcome. An equilibrium is a stable distribution of choices in a population, in the sense that the probability that a player chooses a particular strategy is equal to the frequency of the strategy in the distribution. This equilibrium concept differs fundamentally from the standard game-theoretical analysis and has special properties. For example, a dominated strategy (which is never the best response against any belief about the other player's behavior) might appear with positive probability in the support of the equilibrium when the symmetric game has at least three strategies.

(ii) Rani Spiegler constructed several economic models in which individuals follow an S-1-type procedure (see Spiegler (2011, chapters 6 and 7)). Spiegler (2006a) presents and analyses a model of a market for quacks. The market consists of several healers

whose success rates are identical to the success rate of non-treatment. The healers compete on price. Patients rely on anecdotal evidence regarding the healers' success and also sample the free non-treatment. Spiegler (2006b) presents a model of competition between providers of a service in which each chooses a distribution of price in the range $(-\infty, 1)$. Consumers sample each of the providers and choose the one with the lowest price (which might be negative). The uniqueness of the symmetric Nash equilibrium with expected price of $1/2$ is proven. Increasing the number of competitors leads to a mean-preserving spread in the equilibrium price distribution.

4. Long Interactions and Finite Automata

Story: Two players are involved in a repeated Prisoner's Dilemma interaction. Each of the players chooses one of two modes of behavior each day: C (cooperative) or D (non-cooperative). A player's daily payoff if he chooses the action corresponding to a row while his partner chooses the action associated with a column is given by the corresponding entry in the following matrix:

	C	D
C	3	0
D	4	1

Each period ends with one of four outcomes (C, C) , (C, D) , (D, C) , or (D, D) and the four pairs of payoffs that can be obtained each day are $(3, 3)$, $(0, 4)$, $(4, 0)$, or $(1, 1)$. The infinitely long interaction yields an infinite sequence of outcomes and two streams of payoffs, one for each player. A player cares only about his long-term average payoff.

In the long-term interaction (i.e. a repeated game), each player chooses a strategy, i.e. a plan that specifies whether to play C or D after every possible sequence of his opponent's actions. A strategy can be complex since it specifies an action for an infinite number of contingencies. A player's language in formulating a strategy is a *finite automaton* (machine). This is not an actual machine used to play the strategy but rather an abstract description of the mental process used by a player to determine an action at each point in time after any history he might encounter.

A finite automaton consists of:

- (i) a finite set of states (of mind);
- (ii) an initial state from which the machine starts operating;
- (iii) an output function that indicates, for each state, one action (C or D) which the player will play whenever his machine reaches that state.
- (iv) a transition function which determines the next state of the machine for each state and each possible observation (C or D) of the action taken by the other player.

The two machines are operated in the expected way: The two initial states and the two output functions determine the first-period pair of actions. Each player observes the other player's action and his transition function determines the next state of his machine. The process continues recursively as if each pair of states is the initial one.

Any weighted average with rational coefficients of the four one-shot payoff pairs that gives each player at least 1 (a payoff each player can guarantee by playing D) is a long-term average payoff vector of some pair of strategies that constitutes a Nash equilibrium. In other words, no player can increase his average payoff by using a different strategy. To obtain such a pair of average payoffs (π_1, π_2) , players can play in a cycle of length $n_{D,D} + n_{D,C} + n_{C,D} + n_{C,C}$ so that for all $i, j \in \{C, D\}$ the outcome (i, j) will appear $n_{i,j}$ times in the cycle such that $\frac{n_{D,D}(1,1) + n_{D,C}(4,0) + n_{C,D}(0,4) + n_{C,C}(3,3)}{n_{D,D} + n_{D,C} + n_{C,D} + n_{C,C}} = (\pi_1, \pi_2)$. Players will support this path of play by threatening that any deviation will trigger a permanent switch to non-cooperative behavior.

However, in our story and unlike the standard repeated game model, a player cares not only about his stream of payoffs but also about the complexity of the machine he employs, which is measured by the number of states in the machine (without taking into account the complexity of the transition function). Players view the complexity as a secondary consideration after the desire to increase the average payoff. Thus, a pair of machines will not be stable if one of the players can replace his machine with another that either yields a higher average payoff or the same payoff but lower complexity. In other words, each player has lexicographic preferences with the first priority being to increase average payoff and the second to reduce complexity.

Initially, the two players use a naive strategy - the one-state machine M_C (see Figure 1) which plays C independently of what is observed. This situation is not stable. One of the players modifies his machine to M_D , which plays D independently of what it ob-

serves. Such a deviation improves his payoff each period without increasing complexity.




Figure 1

The two players then adopt the machine M_{grim} . This machine has two states: *friendly* and *punishment*. The initial state is *friendly*. The output function assigns *C* to *friendly* and *D* to *punishment*. If *D* is observed when the machine is at *friendly*, it moves to *punishment*, which is a terminal state. A player who plays against M_{grim} cannot increase his long-term average payoff whatever machine he chooses. However, eventually one of the players notices that he can reduce the complexity of his machine by dropping the state *punishment* without any loss in terms of payoff.

The situation has some happier endings. The following are two examples:

Give and Take: Player 1 chooses the two-state machine M_1 (see Figure 2) with initial state *Take*, in which he plays *D* and moves to the second state *Give* after the other player plays *C*. If M_1 reaches *Give* it plays *C* and moves to *Take* independently of what it observes. Note that “giving” means playing *C* and tolerating the other player playing *D*. “Taking” means playing *D* and expecting the other player to play *C*. Player 2 adopts the machine M_2 which is identical to M_1 except that its initial state is *Give*.




Figure 2

Adopting this pair of machines results in stability. Players expect to alternate between giving and taking. A player’s threat against the other player not giving is to remain in the

Take state and continue taking until the other player gives. As long as player 1 follows M_1 and player 2 follows M_2 , the players will alternate turns between giving and taking and each player obtains an average payoff of 2. No player can reduce the complexity of his machine without losing payoff and no player can increase his average payoff no matter what he does.

The cooperative equilibrium: Each player uses the machine M^* (see Figure 3) which has two states: *Showoff* and *Cooperative*. In *Showoff*, a player plays D and moves to *Cooperative* only if the other player plays D . In *Cooperative*, the machine plays C and moves to *Showoff* only if the other player played D . The machine's initial state is *Showoff*. Thus, players start by proving their ability to punish one another and only then do they move to the cooperative mode.




Figure 3

Each player obtains an average payoff of 3. A player cannot increase his payoff regardless of what he does. Reducing the number of states would cause a payoff loss (the one-state machines M_D and M_C achieve averages of only $5/2$ and 0 , respectively, against M^*).

Discussion: This section is based on Rubinstein (1986) and Abreu and Rubinstein (1988). The bounded rationality component of the model is the complexity of the strategy and the desire of players to minimize it as long as the repeated game payoff is not reduced.

The two equilibria described above demonstrate the logic of two different types of social arrangements. In the first, players alternate between giving and taking; the threat not to give tomorrow deters the other player from not giving today. The second arrangement achieves ideal cooperation but on the way players must demonstrate their ability to punish if necessary.

Abreu and Rubinstein (1988) characterize the Nash equilibria of the infinitely repeated game with finite automata for (i) general one-shot two-player games; (ii) a discounting evaluation criterion; and (iii) any preferences that are increasing in the payoffs and decreasing in complexity. The payoff vectors that can be obtained are those on the “cross” in Figure 4:




Figure 4

A few structural results were proved, including that the two machines must have the same number of states in equilibrium and that there is a one-to-one correspondence between the occurrence of states in the two machines (namely, in equilibrium a machine “knows” the state of the other player’s machine). An improved presentation of the results, following Piccione (1992), appears in Rubinstein (1998, chapter 8).

In a previous paper, Rubinstein (1986) showed that the only equilibria of the machine game with the additional constraint that all states are revisited infinitely often (otherwise they would eventually be dropped) yield either the non-cooperative outcome or cyclical combinations of (C, D) and (D, C) only. Cooperation cannot be obtained in such an equilibrium.

Bibliographic comments:

(i) Neyman (1985) and Ben-Porath (1993) investigate a repeated game with finite automata in which the number of states a player can use in his machine is exogenously bounded. Neyman (1985) shows that in the finite-horizon Prisoner’s Dilemma, cooper-

ation can be achieved in equilibrium by applying the idea that the machines waste their limited resources by following an initial string of actions (“a password to heaven”) that prevents them from identifying the point in time — towards the end of the game — at which it is profitable for them to deviate to the non-cooperative mode of behavior without being punished in the future. Ben-Porath (1993) shows that in an infinitely repeated zero-sum game, if players use the *limit of the means* criterion and are limited in the size of the automaton they can use, then as long as the number of states of one machine is not much larger (in a well-defined sense explained there), the equilibrium payoff vector will be close to the equilibrium payoff of the basic game.

(ii) Examples of other models of repeated games with bounds on the strategies: Lehrer (1988) studies the model of repeated games with bounded recall and Megiddo and Wigder-son (1986) and Chen, Tang and Wang (2017) study the model of repeated games in which players use Turing machines instead of finite automata.

(iii) Eliaz (2003) and Spiegler (2004) study models of repeated games and bargaining where the simplicity considerations were applied to the player’s belief regarding the opponent’s strategy, rather than his own.

5. Agents with different models in mind

Story: Each of two agents, 1 and 2, wishes to purchase a particular service each day. The service is worth 6 to Agent 1 and 4 to Agent 2. Obtaining the service involves a trip with a cost of 3. All involved parties interact daily for a “very long” period of time. The provider sets a daily price. Every morning, without knowing the price charged by P , each agent decides whether to take the trip or not. An agent who comes to P is informed of the price offer and decides whether to buy the service from P or return home. He buys the service if and only if the price does not exceed his valuation. At the end of the period, the price becomes common knowledge to all individuals.

The provider cares only about his average long-run profits. He chooses a sequence of prices. An agent behaves myopically in every period. When he decides whether to make the trip he forms a belief about the price. He makes the trip if and only if his valuation is at least 3 plus the expected price.

An agent bases his belief on the regularities he detects in the sequence of prices. The

ability to detect a regularity is given by a non-negative integer k (referred to as the *order*) which expresses a player's cognitive depth: An agent of order k learns the empirical conditional distribution of the next price given the previous k prices, but cannot detect regularities that require longer histories. Thus, an agent of order 0 only learns the frequency at which different prices appear in the price sequence. An agent of order 1 learns the frequency of each price given the last price. If the real sequence is $(0, 0, 1, 1, 0, 0, 1, 1, \dots)$ an agent of order 0 or 1 will always believe that the next element in the sequence is 0 or 1 with equal probability. An agent of order 2 (or higher) will accurately predict the next element in the sequence since the last two periods fully encode the next element.

Our two agents differ in their ability to recognize patterns. Agent 1 can detect a pattern that determines the next price based on the last three periods ($k_1 = 3$) and Agent 2 is more sophisticated (which aligns with the assumption that his valuation is lower than Agent 1's) and can detect patterns on the basis of the last 4 periods ($k_2 = 4$). Nothing essential would change in the story if we replace these bounds with any other non-negative integers k_1 and k_2 , as long as $k_1 < k_2$. In other words, we assume that the agent who needs the service more has a lower pattern recognition ability.

The provider understands that if he always charges a price higher than 3 then he will not attract any customers; if he constantly charges a price P such that $1 < P \leq 3$, then only Agent 1 will make the trip; and if he always charges a price of 1 or below, then both agents will buy his service but his total profits will be at most 2.

Furthermore, the provider realizes that if both agents have the same cognitive limit, then using random price sequences will not help him since both agents will eventually compare their valuation to the expected price plus 3 and he would not make more than what he would have with a constant price.

Finally, it comes to the provider's attention that he can obtain higher profits if he wisely exploits the differences in the agents' cognitive ability and the correlation between an agent's order and the agent's valuation, namely that the more sophisticated agent has a lower valuation.

He learns about the existence of "De Bruijn sequences of order k " (for every k). To understand this concept, consider the following infinite sequence of prices which has a cycle of length 16 and is an example of a De Bruijn sequence of order 4:

$(5, 5, 5, 5, 1, 5, 5, 1, 1, 5, 1, 5, 1, 1, 1, 1, \dots)$.

An agent with $k = 16$ would, of course, be able to detect a rule that creates this sequence and to correctly predict the next period's price. In fact, any agent with $k \geq 4$ would also be able to. The sequence is built so that all 16 combinations of four prices (1 and 5) appear in the cycle exactly once and thus the last four prices predict the next element in the sequence. On the other hand, the last three values in the sequence predict that it is equally likely that the next price will be 1 or 5. Thus, an agent with $k = 4$ would be able to predict that a low price follows four high prices, that after a sequence of three high prices and one low price always comes a high price, etc. An agent with $k = 3$ assigns probability $1/2$ to each of the two prices after every three-price history.

Thus, Agent 2 will know when the low price will be posted and will approach the provider only on mornings when he expects the price to be 1. Agent 1 is not able to find any useful regularity. After the appearance of any 3 prices, he will believe that there is an equal chance of observing a high or a low price. His expected price is 3, so every morning he approaches the provider and buys the service.

Thus, the provider has found a non-constant sequence, which is complicated enough that agent 1 is left confused and at the same time the sequence is regular enough so that agent 2 is able to predict perfectly. The provider's long-run average profit is $7/2$ per period, which is more than he could obtain using any constant price sequence.

Discussion: This section is based on Piccione and Rubinstein (2003). The bounded rationality element in this model is the limited ability of agents to recognize patterns. Once there is a correlation between this ability and other economic factors, sophisticated manipulators can try to use the correlation as a way to sort agents to their benefit. Conditions for this manipulation are analysed in the paper.

From an economic perspective, this line of research is related to Spence (1973). However, whereas in Spence (1973) the correlation is between two material factors (productivity and cost of acquiring education), here the separation between agents is based on the correlation between a material factor (the willingness to pay) and cognitive ability (to recognize patterns in a time series). Piccione and Rubinstein (2003) also show that the separation can emerge in market equilibrium without the explicit interference of an interested party.

Bibliographic comments:

(i) Rubinstein (1993) is an earlier paper demonstrating the ability of a monopolist to use the correlation between cognitive differences and willingness to pay in order to increase his profits. The modeling of agents' limits in understanding a multi-part price offer makes use of the formal concept of a perceptron.

(ii) Eyster and Piccione (2013) study a model of competitive markets, in which risk-neutral traders trade a one-period bond against an infinitely lived asset in each period. Traders lack structural knowledge of the situation and use diverse and incomplete theories of price determination in order to form statistically correct beliefs about the long-term asset price in the next period. One of the results is that the price of the long-term asset is affected by the diversity in the agents' cognitive levels.

(iii) In a series of papers starting with Jehiel (2005), Philippe Jehiel developed the concept of *analogy-based equilibrium*. The basic idea is related to the way we play games like chess. When planning a move, we evaluate the board positions to which it can lead. In analogy-based equilibrium, those evaluations are correct on average. Players differ in their partition of the set of situations they bundle together, where finer partitions reflect a better understanding of the situation.

(iv) Spiegel (2011, ch. 8) includes a pedagogical exposition of the model and links it to other concepts like Jehiel's.

(v) Rani Spiegel suggests a general framework for studying economic agents who have imperfect understanding of correlation structures and causal relations (Spiegel (2016)). Spiegel (2020) identifies a condition under which the wrong model does not allow an outsider to systematically fool the agent.

References

- Abreu, Dilip and Ariel Rubinstein. 1988. "The Structure of Nash Equilibrium in Repeated Games with Finite Automata". *Econometrica* 56, 1259-1282.
- Ben-Porath, Elchanan. 1993. "Repeated Games with Finite Automata". *Journal of Economic Theory*, 59, 17-32.
- Chen, Lijie, Pingzhong Tang, and Ruosong Wang. 2017. "Bounded Rationality of Restricted Turing Machines". *Proceedings of the AAAI Conference on Artificial Intelligence*, 31, 444-450.
- Choi, Syngjoo, Shachar Kariv, Wieland Müller, and Dan Silverman. 2014. "Who Is (More) Rational?" *American Economic Review*, 104, 1518-50.
- de Clippel, Geoffroy. 2014. "Behavioral Implementation". *American Economic Review*, 104, 2975-3002.
- Eliaz, Kfir. 2003. "Nash Equilibrium When Players Account for the Complexity of their Forecasts". *Games and Economic Behavior*, 44, 286-310.
- Eyster, Erik and Michele Piccione. 2013. "An Approach to Asset Pricing Under Incomplete and Diverse Perceptions", *Econometrica*, 81, 1483-1506.
- Glazer, Jacob and Ariel Rubinstein. 1996. "An Extensive Game as a Guide for Solving a Normal Game", *Journal of Economic Theory*, 70, 32-42.
- Glazer, Jacob and Ariel Rubinstein. 2012. "A Model of Persuasion with Boundedly Rational Agents". *Journal of Political Economy*, 120, 1057-1082.
- Glazer, Jacob and Ariel Rubinstein. 2014. "Complex Questionnaires". *Econometrica*, 82, 1529-1541.
- Jehiel, Philippe. 2005. "Analogy-Based Expectation Equilibrium". *Journal of Economic Theory*, 123, 81-104.
- Lehrer, Ehud. 1988. "Repeated Games with Stationary Bounded Recall Strategies". *Journal of Economic Theory*, 46, 130-144.
- Li, Shengwu. 2017. "Obviously Strategy-Proof Mechanisms". *American Economic Review*, 107, 3257-87.
- Megiddo, Nimrod and Avi Wigderson. 1986. "On Play by Means of Computing Machines". In *Theoretical Aspects of Reasoning About Knowledge: Proceedings of the 1986 Conference*, 259-274.

Neyman, Abraham. 1985. "Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoners' Dilemma". *Economics Letters*, 19, 227-229.

Osborne, Martin and Ariel Rubinstein. 1998. "Games with Procedurally Rational Players". *American Economic Review*, 88, 834-847.

Osborne, Martin and Ariel Rubinstein. 2003. "Sampling Equilibrium with an Application to Strategic Voting". *Games and Economic Behavior*, 45, 434-441.

Piccione, Michele. 1992. "Finite Automata Equilibria with Discounting." *Journal of Economic Theory*, 56, 180-193.

Piccione, Michele and Ariel Rubinstein. 2003. "Modeling the Economic Interaction of Agents with Diverse Abilities to Recognize Equilibrium Patterns". *Journal of the European Economic Association*, 1, 212-223.

Rubinstein, Ariel. 1986. "Finite Automata Play the Repeated Prisoner's Dilemma". *Journal of Economic Theory*, 39, 83-96.

Rubinstein, Ariel. 1993. "On Price Recognition and Computational Complexity in a Monopolistic Model". *Journal of Political Economy*, 101, 473-484.

Rubinstein, Ariel. 1998. *Modeling Bounded Rationality*. MIT Press.

Rubinstein, Ariel. 2012. *Economic Fables*. Open Book Publishers.

Spence, Michael. 1973. "Job Market Signaling". *Quarterly Journal of Economics*, 87, 355-374.

Spiegler, Ran. 2004. "Simplicity of Beliefs and Delay Tactics in a Concession Game". *Games and Economic Behavior*, 47, 200-220.

Spiegler, Ran. 2006a. "The Market for Quacks". *Review of Economic Studies*, 73, 1113-1131.

Spiegler, Ran. 2006b. "Competition over Agents with Boundedly Rational Expectations". *Theoretical Economics*, 1, 207-231.

Spiegler, Ran. 2011. *Bounded Rationality and Industrial Organization*. Oxford University Press.

Spiegler, Ran. 2016. "Bayesian Networks and Boundedly Rational Expectations". *Quarterly Journal of Economics*, 131, 1243-1290.

Spiegler, Ran. 2020. "Can Agents with Causal Misperceptions be Systematically Fooled?" *Journal of the European Economic Association*, 18, 583-617.