

---

# Modeling Bounded Rationality of Agents During Interactions

---

**Qing Guo**

Department of Computer Science  
University of Illinois at Chicago  
Chicago, IL 60607  
qguo@cs.uic.edu

**Piotr Gmytrasiewicz**

Department of Computer Science  
University of Illinois at Chicago  
Chicago, IL 60607  
piotr@cs.uic.edu

## Abstract

Frequently, it is advantageous for an agent to model other agents in order to predict their behavior during an interaction. Modeling others as rational has a long tradition in AI and game theory, but modeling other agents' departures from rationality is difficult and controversial. This paper proposes that bounded rationality be modeled as errors the agent being modeled is making while deciding on actions. We are motivated by the work on quantal response equilibria in behavioral game theory which uses Nash equilibria as the solution concept. In contrast, we use decision-theoretic maximization of expected utility and apply the models within the framework of interactive POMDPs. Quantal response assumes that a decision maker is rational, i.e., is maximizing his expected utility, but only approximately so, with an error rate characterized by a single error parameter. Another agent's error rate may be unknown and needs to be estimated during an interaction. We show that the error rate of the quantal response can be estimated using Bayesian update of a suitable conjugate prior, and that it has a sufficient statistic under strong simplifying assumptions. However, if the simplifying assumptions are relaxed, the quantal response does not admit a finite sufficient statistic and a more complex update is needed. This confirms the difficulty of using simple models of bounded rationality in general settings.

## 1 Introduction

In AI, an agent's (perfect) rationality is defined as the agent's ability to execute actions that, at every instant, maximize the agent's expected utility, given the information it has acquired from the environment [13]. Let us note two aspects of this definition. First, the fact that the acquired information may be limited does not preclude perfect rationality. In other words, an agent may have very limited information but still be perfectly rational. Second, the above definition does not specify any particular procedure an agent is to use to decide which action to execute. Further, the definition is completely independent of any details of the implementation of any such decision making procedure.

The notion of bounded rationality received a lot of attention in economics and psychology. Simon [14] coined the term and suggested it as an alternative to rationality. Simon pointed out that perfectly rational decision making is often difficult in practice due to limited cognitive and/or computational resources. He proposed that humans are *satisficers*, as opposed to perfect optimizers, and that they use heuristics to make decisions, rather than optimization rules. Gigerenzer [7, 6] argued that simple heuristics could actually lead to better decisions than theoretically optimal procedures. The use of heuristics was also studied by Kahneman [9], who proposed his own alternative to perfect rationality called prospect theory. Rubinstein [12] proposed that one needs to model an agent's decision-making procedures explicitly in order to model the agent's bounded rationality adequately.

This paper builds on an approach to modeling bounded rationality called quantal response [2, 10, 11]. It is a simple model which uses a single error parameter. Quantal response is simple in that it does not attempt to model the procedures, and their possible limitations, the agent may use to decide on action. The great advantage of this model is that, first, there exist a myriad of procedural mechanisms by which perfect rationality could be implemented, heuristics which could be used, and possible ways in which any of these could have its functionality limited by the specific computational or cognitive architecture of the agent in question. Second, none of these implementation details and architectural limitations are observable by the external observer who is doing the modeling. In other words, quantal response abstracts away the unobservable parameters specific to implementation and treats them as *noise* which produces non-systematic departures from perfect rationality.

The models of bounded rational agents we are constructing are intended to be used within interactive partially observable Markov decision processes (I-POMDPs) [8]. In I-POMDPs, agents maintain beliefs about the state of the world, as in classical POMDPs, and about the other agents with whom they interact. These beliefs are probability distributions spanned on the space of all possible states of the environment and all possible models of the other agents. Agents best respond to their beliefs using decision-theoretic expected utility maximization, like in classical POMDPs. Intentional models of other agents specify their preferences (payoff functions), action sets, and beliefs, and are analogous to POMDP-like specifications of the other agents' decision-making. The probabilities are assigned to various models since the internals of the other agents' specifications are not directly observable. In other words, an agent may not know before hand, say, the payoff function of another agent, but may be able to probabilistically infer their likely payoffs by observing their actions.

To make room for bounded rationality of the other agents, we define a notion of approximately intentional agent model. It is analogous to perfectly rational agent model but with a noise factor inversely proportional to an error parameter,  $\lambda$ . According to quantal response [2, 10, 11], probabilities of actions are given by the logit function of the actions' expected utilities. Thus actions that are suboptimal are possible, but their probabilities increase with their expected utilities.

Quantal response specifies the probabilities of an agent's actions given their expected utilities and the agent's error parameter,  $\lambda$ . An additional complication is that an agent's error parameter is not directly observable. Instead, it must be inferred based on the agent's observed behavior. We take a Bayesian approach to this and propose that the modeling agent maintain a probability distribution over possible values of  $\lambda$  for the modeled agent, and that this probability be updated when new actions are observed. Intuitively, if an agent is observed acting rationally then, over time the error rate attributed to this agent should decrease (and since  $\lambda$  is an inverse error, larger values of  $\lambda$  should become more likely). If, on the other hand, the modeled agent is frequently observed acting in ways that depart from perfect rationality, then the error rate attributed to it should increase (and smaller values of  $\lambda$  should become more likely).

Below we show how the update of the error parameter modeling bounded rationality of another agent can be performed. We also show that in simple special cases, when the interaction is *episodic*, the error rate admits a sufficient statistic. We then derive a distribution over  $\lambda$  that is a member of a family of conjugate priors. That means that the update of the distribution over  $\lambda$  is particularly simple and that it results in another distribution in the same family of parametrized distributions. We further show that if the simplifying assumptions are relaxed, then there is no sufficient statistic and no conjugate prior over  $\lambda$ . In these cases we derive the update of the error parameter within the general, and more complex, belief update in I-POMDPs.

## 2 Logit Quantal Response

For simplicity, we assume that a modeling agent, called  $i$ , is considering the behavior of one other agent,  $j$ . The logit quantal response is defined as follows [2, 10, 11]:

$$P(a_j) = \frac{e^{\lambda u_{a_j}}}{\sum_{l=1}^m e^{\lambda u_{a_l}}}, \quad (1)$$

where  $\{a_l : l = 1, 2, 3, \dots, m\}$  is a set of all possible actions of the agent.  $P(a_j)$  is the probability of the agent  $j$  taking the action  $a_j$ .  $u_{a_j} \in \mathbb{R}$  is the expected utility of action  $a_j$  to agent  $j$  and  $\lambda \geq 0$  is the (inverse) error rate of the agent  $j$ .  $\lambda$  represents how rational agent  $j$  is: greater  $\lambda$  makes it more

likely that  $j$  takes actions which have higher utilities. When  $\lambda \rightarrow +\infty$ ,  $P(a_j) = 1$  for the action which has the highest expected utility<sup>1</sup> and  $P(a_j) = 0$  for all other actions. This means agent  $j$  is perfectly rational because he always chooses an action with the best expected utility. When  $\lambda = 0$ ,  $P(a_j) = 1/m$ ,  $\forall j = 1, 2, 3, \dots, m$ , which means agent  $j$  chooses actions at random.

It is likely that the error rate  $\lambda$  of agent  $j$  is not directly observable to agent  $i$ . Bayesian approach allows agent  $i$  to learn this rate during interactions. To do this agent  $i$  needs a prior distribution,  $f(\lambda)$ , which represents  $i$ 's current knowledge about agent  $j$ 's error rate, and to observe agent  $j$ 's action,  $a_j$  at the current step. The updated distribution is:

$$f(\lambda|a_j) = \frac{P(a_j|\lambda)f(\lambda)}{\int_0^\infty P(a_j|\lambda')f(\lambda') d\lambda'}. \quad (2)$$

Using the above formula, agent  $i$  can maintain his knowledge about agent  $j$ 's bounded rationality by repeatedly updating  $f(\lambda)$  during interaction.

Formula (2) may not be easy to apply because after updating the  $f(\lambda)$  several times, it becomes more and more complicated. To overcome this it is convenient to look for a conjugate prior family. In Bayesian probability, if the posterior distribution is in the same family as the prior distribution, then this prior is called a *conjugate prior* [3, 4]. Conjugate priors are convenient because they make the updating process tractable; one just needs to update the parameters of the conjugate prior distribution (hyperparameters) to realize the Bayesian update.

### 3 Static Episodic Environments with Perfect Observability

In this section we consider the simplest case, when agent  $j$ 's expected utilities  $u_{a_l}$  for all actions are *known* to agent  $i$  and remain the same during the interaction. In other words, agent  $j$  is not updating his beliefs since the environment is static and episodic [13] and  $i$  is observing  $j$  acting in the same decision-making situation repeatedly. The derivation below follows techniques in [3, 4].

Consider the following family of distributions over  $\lambda$ :

$$f(\lambda; u, n) = \frac{e^{\lambda u} / (\sum_{l=1}^m e^{\lambda u_{a_l}})^n}{\int_0^\infty e^{\lambda' u} / (\sum_{l=1}^m e^{\lambda' u_{a_l}})^n d\lambda'}, \quad (3)$$

where  $n$  and  $u$  are hyperparameters. Here  $n$  is a natural number including zero, and  $u$  is restricted by following:  $u < n \max_l u_{a_l}$ . One can verify (3) is a probability density function since  $\int_0^\infty f(\lambda; u, n) d\lambda = 1$ , given the restrictions over  $n$  and  $u$ .

**Proposition 1:** The family of distributions  $f(\lambda; u, n)$  in (3) is a conjugate family of distributions over  $\lambda$  in static episodic environments with known utilities of actions.

The proof of Proposition 1 establishes how to update the hyperparameters of our conjugate prior after observing that agent  $j$  executed his action  $a_j$ , with expected utility  $u_{a_j}$ :

$$f(\lambda; u, n) \xrightarrow{a_j} f(\lambda; u + u_{a_j}, n + 1). \quad (4)$$

Note that the integral in the denominator of  $f(\lambda; u, n)$  does not always have an analytical solution, so we have to use numerical methods to calculate its value.

One can verify that once there is a valid prior, all the posteriors are always valid. The question also arises as to what is an appropriate prior agent  $i$  should choose before any observations. Often one looks for an uninformed prior. In our case  $f(\lambda; -\epsilon, 0)$ , where  $\epsilon > 0$  is a small positive value, is such an uninformed prior; it is almost flat over the positive real values of  $\lambda$ , as we show in the example below.

---

<sup>1</sup>If there are many, say  $h$ , optimal actions with the same expected utilities, then  $P(a_j) = 1/h$  for each of them.

## 4 Example




Figure 1: Example conjugate prior:  $f(\lambda; -0.1, 0)$

Let us assume that agent  $j$  chooses from among three ( $m = 3$ ) actions, with following expected utilities:  $u_{a_1} = 0, u_{a_2} = 2, u_{a_3} = 10$ . As we mentioned we assume that the expected utilities of agent  $j$  are known to agent  $i$ , and that they do not change. Let the prior be  $f(\lambda; -0.1, 0)$ . Let us first compute the expected value of the error parameter  $i$  attributes to  $j$  under this distribution:  $E(\Lambda) = \int_0^\infty \lambda f(\lambda) d\lambda = 10.0$ . Using the formula of total probability for each action of  $j$  we get:  $P(a_j) = \int_0^\infty P(a_j|\lambda) f(\lambda) d\lambda$ . Thus the prior probabilities  $i$  attributes to each of  $j$ 's actions are:  $P(a_1) = 0.00524, P(a_2) = 0.00699, P(a_3) = 0.98777$ . Figure 1 shows the initial prior. Note that this uninformative prior assigns relatively high probabilistic weight to high values of  $\lambda$  and hence high degree of  $j$ 's rationality.

$\lambda$ 's Distribution	$E(\Lambda)$	$P(a_1)$	$P(a_2)$	$P(a_3)$
$f(\lambda; -0.1, 0)$	10.0000	0.00524	0.00699	0.98777
$f(\lambda; 29.9, 3)$	10.2477	0.00138	0.00229	0.99633
$f(\lambda; 299.9, 30)$	10.5309	0.00010	0.00029	0.99961
$f(\lambda; 11.9, 3)$	0.1069	0.20663	0.24147	0.55190
$f(\lambda; 119.9, 30)$	0.0328	0.28959	0.30785	0.40256

Table 1: Probabilities of agent  $j$ 's actions derived from various distributions over error parameter  $\lambda$ .




Figure 2:  $f(\lambda; 29.9, 3)$ , updated after observing three rational actions.

Assume agent  $j$  acts rationally and always chooses his best action,  $a_3$ . Then Figure 2 and Figure 3 show the posterior after three observations ( $f(\lambda; 29.9, 3)$ ) and after 30 observations ( $f(\lambda; 299.9, 30)$ ) of  $j$ 's action  $a_3$ . We can see that higher values of  $\lambda$  become more likely if the agent always chooses the action with the best utility. We can also compute the probabilities of the three actions under these two posteriors, which are shown in Table 1.

Now let us assume that agent  $j$  behaves randomly. Within the first three actions, he chooses each of his actions  $a_1, a_2$  and  $a_3$  once. The updated distribution over  $i$ 's error parameter is then  $f(\lambda; 11.9, 3)$ ,




Figure 3:  $f(\lambda; 299.9, 30)$ , updated after observing 30 rational actions.




Figure 4:  $f(\lambda; 11.9, 3)$ , updated after observing three random actions.

which is shown in Figure 4. Further, if within  $j$ 's 30 actions he chooses  $a_1$  for ten times,  $a_2$  for ten times, and  $a_3$  for ten times; then the posterior is  $f(\lambda; 119.9, 30)$ , which is shown in Figure 5. The results are intuitive. Thus, if agent  $j$  behaves randomly, lower values of  $\lambda$ , indicating stronger departure from perfect rationality, become more likely. Probabilities of the three actions under these two posteriors are also shown in Table 1.

## 5 Sequential Dynamic Environment with Perfect Observability of Finite Types

In this section, we extend our approach to more complex case of dynamic sequential environment. Again, we assume that expected utilities of  $j$ 's actions are known to  $i$ , but now, since agent  $j$  may be updating his beliefs, the expected utilities of his actions do not remain constant but can take a finite number of values. We refer to each of the beliefs of agent  $j$ , together with his payoff function and other elements of his POMDP, as  $j$ 's type,  $\theta_j$ . Thus, the set of possible types of agent  $j$ ,  $\Theta_j$ , has  $K$  possible elements  $1, 2, \dots, K$ . We denote  $U(a_j|\theta_j = k) = u_{a_j,k}$ , where  $k = 1, 2, \dots, K$  and assume that index  $k$  is observable (or computable) by agent  $i$ . Then the logit quantal response (1) for the probability of agent  $j$  taking action  $a_j$  given his  $k$ th type is:

$$P(a_j|k, \lambda) = \frac{e^{\lambda u_{a_j,k}}}{\sum_{l=1}^m e^{\lambda u_{a_l,k}}}. \quad (5)$$

Now Bayesian update, analogous to equation (2), becomes:

$$f(\lambda|a_j, k) = \frac{P(a_j|k, \lambda)f(\lambda)}{\int_0^\infty P(a_j|k, \lambda')f(\lambda')d\lambda'}. \quad (6)$$




Figure 5:  $f(\lambda; 119.9, 30)$ , updated after observing 30 random actions.

We now have a proposition analogous to Proposition 1 in Section 3. Consider the following family of distributions:

$$f(\lambda; u, n_1, n_2, \dots, n_K) = \frac{e^{\lambda u} / \prod_{k=1}^K (\sum_{l=1}^m e^{\lambda u_{a_l, k}})^{n_k}}{\int_0^\infty e^{\lambda' u} / \prod_{k=1}^K (\sum_{l=1}^m e^{\lambda' u_{a_l, k}})^{n_k} d\lambda'}, \quad (7)$$

where  $n_k = 0, 1, \dots, \forall k = 1, \dots, K$ ;  $u < \sum_{k=1}^K (n_k \max_l u_{a_l, k})$ . One can verify that (7) is a valid probability density function since integral of the denominator converges if and only if  $u < \sum_{k=1}^K (n_k \max_l u_{a_l, k})$ .

**Proposition 2:** The family of distributions in (7),  $f(\lambda; u, n_1, n_2, \dots, n_K)$  is a conjugate family of distributions over  $\lambda$  in a sequential dynamic environment with perfect observability of finite number of types.

Similarly to the simpler case of Proposition 1, the proof of Proposition 2 establishes the update of the hyperparameters of the conjugate prior based on the observed action,  $a_j$ , with expected utility  $u_{a_j, k}$ :

$$f(\lambda; u, n_1, n_2, \dots, n_K) \xrightarrow{a_j, k} f(\lambda; u + u_{a_j, k}, n_1, n_2, \dots, n_{k-1}, n_k + 1, n_{k+1}, \dots, n_K). \quad (8)$$

Similarly to Section 3, once there is a valid prior, e.g.  $f(\lambda; u, n_1, n_2, \dots, n_K)$ , all the posteriors are always valid. An uninformative prior agent  $i$  can choose before observing any of  $j$ 's actions can be  $f(\lambda; -\epsilon, 0, 0, \dots, 0)$ . Then after any number of observations the current  $u$  is the accumulated utility of all actions the agent has taken minus  $\epsilon$ , and current  $n_k$  is the counter of occurrence of the  $k$ th type.

## 6 Sequential Dynamic Environments with Perfect Observability of Continuous Types

Let us consider an even more general case, in which the expected utilities  $u_{a_l}$  are not limited to a finite number of values but can lie in some interval or even on the real line:

$$P(a_j | \mathbf{u}, \lambda) = \frac{e^{\lambda u_{a_j}}}{\sum_{l=1}^m e^{\lambda u_{a_l}}}, \quad (9)$$

where  $u_l < u_{a_l} < u_l', l = 1, 2, \dots, m$ ,  $u_l \geq -\infty$  and  $u_l' \leq \infty$  are lower and upper bounds of the expected utilities  $u_{a_l}$ , and where  $\mathbf{u}$  is a vector of expected utilities of all  $m$  actions,  $\mathbf{u} = (u_{a_1}, u_{a_2}, \dots, u_{a_m})$ . Again assume  $u_{a_l}$  are known to agent  $i$ , and he observes agent  $j$ 's action  $a_j$ .

Similarly to Section 5, the Bayesian update equation with continuous types is

$$f(\lambda | a_j, \mathbf{u}) = \frac{P(a_j | \mathbf{u}, \lambda) f(\lambda)}{\int_0^\infty P(a_j | \mathbf{u}, \lambda') f(\lambda') d\lambda'}. \quad (10)$$

If we want to update the distribution of  $\lambda$  it would be convenient to find a conjugate prior of (9) for  $\lambda$ . However, forming a conjugate prior in this case is not easy, and may be impossible. The reason is that the construction of conjugate prior distributions [3, 4] is based on the existence of sufficient statistics of fixed dimension for the given likelihood function (formula (9) in this case). However, under very weak conditions, the existence of fixed dimensional sufficient statistic restricts the likelihood function to the exponential family of functions [1, 5]. Unfortunately, (9) does not belong to the exponential family with continuous utilities  $\mathbf{u}$  when  $m \geq 2$ .

In other words, in this case, there is no known way of deriving a family of conjugate priors. Two ways of circumventing this difficulty present themselves. First is to discretize  $\mathbf{u}$  and approximate it by fitting its values into a finite number of types. The second one is to give up on conjugate priors altogether and use numerical approximation to update  $\lambda$ . We pursue this last alternative in the section below, but in the more general case when the types of agent  $j$  are not observable.

## 7 Sequential Dynamic Environments with Unobservable Types: I-POMDPs

It turns out that the logit quantal response can be included in interactive POMDPs (I-POMDPs) [8] to model bounded rationality of agents without assuming that their types (i.e., payoff functions and beliefs) are known to other agents. As we mentioned, we do this by allowing agent  $i$  to define approximately intentional models of agent  $j$  by including an error rate,  $\lambda$ , into  $j$ 's type. We define agent  $i$ 's I-POMDP as:

$$I\text{-POMDP}_i = \langle IS_i, A, T_i, \Omega_i, O_i, R_i \rangle, \quad (11)$$

where:

- $IS_i = S \times \Theta_j^\lambda$  is a set of interactive states, where  $S$  is the set of physical states and  $\Theta_j^\lambda$  is a set of possible approximately intentional models of agent  $j$ . An approximately intentional model of  $j$  is defined as  $\theta_j^\lambda = \langle b_j, \hat{\theta}_j^\lambda \rangle$ , where  $b_j$  is agent  $j$ 's belief state, and  $\hat{\theta}_j^\lambda = \langle \hat{\theta}_j, \lambda_j \rangle$ . In  $\hat{\theta}_j^\lambda$ ,  $\hat{\theta}_j = \langle A, \Omega_j, T_j, O_j, R_j, OC_j \rangle$  is  $j$ 's frame, and  $\lambda_j$  is  $j$ 's (inverse) error rate, as before. Agent  $j$ 's belief  $b_j$  is a probability distribution over the physical states and agent  $i$ 's approximately intentional models, namely  $b_j \in \Delta(S \times \Theta_i^\lambda)$ .
- $A = A_i \times A_j$  is the set of all joint actions of the agents.
- $T_i : S \times A \times S \rightarrow [0, 1]$  is agent  $i$ 's transition function.
- $\Omega_i$  is agent  $i$ 's possible observations.
- $O_i : S \times A \times \Omega_i \rightarrow [0, 1]$  is agent  $i$ 's observation function.
- $R_i : IS_i \times A \rightarrow \mathbb{R}$  is agent  $i$ 's reward function.

If agent  $i$  is modeling agent  $j$  as an approximately intentional I-POMDP agent who behaves rational with errors represented by logit quantal response, then agent  $i$  updates his beliefs within I-POMDPs as:

$$\begin{aligned} b_i^t(is^t) = & \beta \sum_{is^{t-1}: \hat{\theta}_j^{\lambda, t-1} = \hat{\theta}_j^{\lambda, t}} b_i^{t-1}(is^{t-1}) \sum_{a_j^{t-1}} Pr(a_j^{t-1} | \theta_j^{\lambda, t-1}) O_i(s^t, a^{t-1}, o_i^t) \\ & \times \sum_{o_j^t} \tau_{\theta_j^{\lambda, t}}(b_j^{t-1}, a_j^{t-1}, o_j^t, b_j^t) O_j(s^t, a^{t-1}, o_j^t) T_i(s^{t-1}, a^{t-1}, s^t), \end{aligned} \quad (12)$$

in which  $Pr(a_j^{t-1} | \theta_j^{\lambda, t-1})$  is given by logit quantal response:

$$Pr(a_j^{t-1} | \theta_j^{\lambda, t-1}) = \frac{e^{\lambda_j^{t-1} U(a_j | \theta_j^{t-1})}}{\sum_{l=1}^m e^{\lambda_j^{t-1} U(a_l | \theta_j^{t-1})}}. \quad (13)$$

Given  $j$ 's type  $\theta_j$ , the expected utilities of each action is

$$U(a_j|\theta_j) = \sum_{is} ER_j(is, a_j)b_j(is) + \gamma \sum_{o_j \in \Omega_j} Pr(o_j|a_j, b_j)U(\langle SE_{\theta_j}(b_j, a_j, o_j), \hat{\theta}_j \rangle), \quad (14)$$

where  $a_j \in A_j$  ( $A_j$  is  $j$ 's action set).

## 8 Conclusion

In this paper we postulated that bounded rationality of agents be modeled as noise, or error rate, that perturbs their rational action selection. Since error rates of other agents are not directly observable we presented ways to learn these parameters during interactions. The learning uses Bayesian update, for which it is convenient to use a family of conjugate priors over the possible values of the error rate. We found the conjugate priors of logit quantal response functions for static and episodic environments, and for sequential dynamic environments with finite number of observable types. The existence of conjugate priors under these assumptions makes the task of learning another agent's error rate simple and tractable. However, we have also shown that if the space of types of the modeled agent is continuous, then the quantal response likelihood does not satisfy the precondition needed for construction of conjugate priors over the error rates. Discretizing their utilities to make continuous types fit into finite pre-specified types can be a way of solving this difficulty. Another method is to abandon the search for conjugate priors and use general updates. These updates, derivable from the I-POMDP framework, do not assume that other agents' payoff functions and beliefs are known, but are quite difficult to compute in realistic cases.

## References

- [1] O. Barndorff-Nielsen and K. Pedersen. Sufficient data reduction and exponential families. *Math. Scand.*, 22:197–202, 1968.
- [2] C. F. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, 2003.
- [3] M. H. DeGroot. *Optimal Statistical Decisions (Wiley Classics Library)*. Wiley-Interscience, April 2004.
- [4] D. Fink. A Compendium of Conjugate Priors. Technical report, 1997.
- [5] D. A. S. Fraser. On sufficiency and the exponential family. *Journal of the Royal Statistical Society. Series B (Methodological)*, 25(1):115–123, 1963.
- [6] G. Gigerenzer, editor. *Adaptive thinking: Rationality in the real world*. Oxford University Press, 2000.
- [7] G. Gigerenzer and D. G. Goldstein. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4):650–669, 1996.
- [8] P. Gmytrasiewicz and P. Doshi. A framework for sequential planning in multiagent settings. *Journal of Artificial Intelligence Research*, 24:49–79, 2005. <http://jair.org/contents/v24.html>.
- [9] D. Kahneman, P. Slovic, and A. Tversky. *Judgment under uncertainty: Heuristics and biases*. Cambridge University Press, 1982.
- [10] R. McKelvey and T. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10:6–38, 1995.
- [11] R. McKelvey and T. Palfrey. Quantal response equilibria for extensive form games. *Experimental Economics*, 1:9–41, 1998.
- [12] A. Rubinstein. *Modeling Bounded Rationality*. MIT Press, 1998.
- [13] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach (Third Edition)*. Prentice Hall, 2010.
- [14] H. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69:99–118, 1955.