

A Formal Separation Between Strategic and Nonstrategic Behavior

James R. Wright^{a,*}, Kevin Leyton-Brown^b

^aUniversity of Alberta, Canada

^bUniversity of British Columbia, Canada

Abstract

It is common to make a distinction between “strategic” behavior and other forms of intentional but “non-strategic” behavior: typically, that strategic agents model other agents while nonstrategic agents do not. However, a crisp boundary between these concepts has proven elusive. This problem is pervasive throughout the game theoretic literature on bounded rationality. It is particularly critical in parts of the behavioral game theory literature that make an explicit distinction between the behavior of “nonstrategic” level-0 agents and “strategic” higher-level agents (e.g., the level- k and cognitive hierarchy models). The literature gives no clear guidance on how the rationality of nonstrategic agents must be bounded, instead typically just singling out specific decision rules and informally asserting them to be nonstrategic (e.g., truthfully revealing private information; randomizing uniformly). In this work, we propose a new, formal characterization of nonstrategic behavior. Our main contribution is to show that it satisfies two properties: (1) it is general enough to capture all purportedly “nonstrategic” decision rules of which we are aware; (2) behavior that obeys our characterization is distinct from strategic behavior in a precise sense.

1. Introduction

Our focus in this paper is on *game-theoretic environments*, characterized by multiple independent agents with independent goals, each of whose payoffs depend upon more than just their own actions. How should agents be assumed to behave in such environments? The classical answer is that agents should be modeled as *fully rational*. This is an extremely demanding standard. For example, when a fully rational agent i 's payoffs depend upon the actions of some other self-interested agent j , i must form probabilistic beliefs about j 's actions. If i believes j to be fully rational as well, j 's actions are thus responses to beliefs about i , and so i must also hold beliefs about j 's beliefs about i , and about j 's beliefs about i 's beliefs about j , and so on. Although such assumptions may seem baroque, they in fact undergird the predictions of classical game theory, which assumes all agents to be fully rational; e.g., one way of defining the Nash equilibrium is as such a set of beliefs in which every agent is correct in its beliefs about the others. The behavior of fully rational agents is commonly said to be *strategic*. Indeed, in the early days of game theory, the term “strategic” was used as a synonym for full rationality [e.g., 5, 23].

*Corresponding author

At the other extreme, we might assume that agents are entirely irrational and do not model other agents at all, following some fixed rule like truthfully revealing all private information or uniformly randomizing across all available actions. The behavior of such agents is commonly said to be *nonstrategic*.

Things get muddier in between these extremes. Human players are clearly not fully rational; e.g., nobody knows the Nash equilibrium strategy for chess. However, at least some of us surely do reason about the behavior and beliefs of other agents with whom we interact. The literature generally also calls such “boundedly rational” behavior strategic, even when the behavior is inconsistent with full rationality; intuitively, the dividing line is generally taken to be the question of whether agents model other agents and their incentives when deciding how to act. More formally, the term “strategic” is generally used¹ to describe agents who act to maximize their own utilities based on explicit probabilistic beliefs about the actions of other agents [24, 18, 3, 17, 12, 13, 15], and the term “nonstrategic” is generally used to describe agents who follow some fixed, known decision rule, often truth-telling [26, 2, 18, 17, 12, 15]. The key theme of this paper is that the task of making this idea precise is both underexplored in the literature and tricky in practice.

The literature has taken two main approaches to conceptualizing bounded rationality. The first views boundedness as error. For example, quantal response notions, especially quantal response equilibrium, assume that agents make an error by choosing a less-than-optimal action with increasing probability as the payoffs from taking the suboptimal action approach the payoffs from taking the optimal action [19, 30]. An ϵ -equilibrium is an equilibrium in which every agent comes within ϵ of best responding to the others [28]. More generally, a common approach to modeling less than fully rational agents is to mix in uniform noise representing the probability of making an error [e.g., 9].

The other main approach to conceptualizing bounded rationality views boundedness as a structural property of an agent’s reasoning process. Some work distinguishes between distinct forms of reasoning. For example, Sandholm and Lesser [26] distinguish between strategic threshold strategies and Nash threshold strategies; work on the mean field equilibrium solution concept distinguishes between “cognizant” strategies, which track the current state of other agents, and “oblivious” strategies, which do not [1]. Other work considers finer gradations between the levels of reasoning that different agents are able to perform, as in bounded rationalizability [5, 23] and work in program equilibrium, which often models boundedness by a bounded number of states in a finite state automaton [14, 25].

Another example in this latter vein serves as the most concrete application of both strategic and non-strategic descriptions of behavior within a single model of which we are aware: the *iterative models* of behavioral game theory. These models thus serve as a running example throughout the paper and as one application area in which our work has immediate implications. Overall, the field of *behavioral game theory* aims to develop models that more accurately describe human behavior in game-theoretic environments, based

¹Of course, the term “strategic” is not used entirely consistently in the literature. For example, some work refers to strategic agents as those that pay any attention to the effects of their own actions when reasoning about how to maximize their own utility [31, 2, 4]. Under this meaning for the term, there exists a meaningful distinction between nonstrategic and strategic agents even in single-agent environments: the former are myopic about the effects of their own actions.

on both field data and experimental data [e.g., 7]. Iterative models are one prominent class of models from this literature; they include the level- k [21, 9, 10], cognitive hierarchy [6], and quantal cognitive hierarchy models [33]. In all of these models, each agent has a non-negative integer *level* representing the degree of strategic reasoning (i.e., modeling of recursive beliefs) of which the agent is capable. Level-0 agents are *nonstrategic*—they do not model other agents’ beliefs or actions at all; level-1 agents model level-0 agents’ actions; level-2 agents model the beliefs and actions of level-1 agents; and so forth.

So far so good; the catch comes when we start getting fancy with the definition of level-0 agents. In most work in the literature, the issue does not come up; level-0 behavior is defined simply as uniform randomization. However, we showed in recent work that model performance can be substantially improved by allowing for richer level-0 specifications [32, 16, 34]. This raises the question of how rich these specifications should be allowed to become, before level-0 stops being plausible as a description of nonstrategic behavior. In our past work we followed the AI tradition described earlier, saying that a strategic agent’s behavior can depend on an explicit, probabilistic belief about the other agents’ actions while a non-strategic agent’s behavior must not. But this argument has a crucial weakness: just because a proposed level-0 behavior can be written without reference to beliefs about other agents’ strategies, we cannot conclude that there does not exist another, equivalent way of writing it that does depend on such beliefs. Things get even worse if we aspire to learn the level-0 specification directly from data, effectively optimizing over a space of specifications, as we do in our most recent work [16]: the task now becomes reassuring a skeptic that no point in this space corresponds to behavior that could somehow be rewritten in strategic terms.

This paper establishes a firm and constructive dividing line between nonstrategic and strategic behaviors. Specifically, it characterizes a broad family of “nonstrategic” decision rules and shows that they deserve the name: i.e., that no rule in this class can represent strategic reasoning. In the sense of the literature on bounded rationality just described, our proposed characterization is a structural notion: it restricts the information that agents are permitted to use by restricting them to summarize all outcomes into a single number before performing their reasoning. In what follows, we introduce notation and important background in Section 2. In particular, we introduce the concept of *behavioral models* in Section 2.3: functions that map from an arbitrary game to a probability distribution over a single agent’s actions in that game. We define strategic behavioral models in Section 3; these are behavioral models that require agents to take account of the incentives of the other agents and to be self interested, both in a precise sense. In Section 4, we prove that all of the existing solution concepts described in Section 2 are strategic. Section 5 then defines *elementary* behavioral models—agent behaviors that can be computed in terms of a matrix of real-valued *potentials* that score each outcome in the game independently—and shows how a wide range of “level-0” behaviors from the literature can be encoded as elementary models. In Section 6, we prove our main result: that minimally self-interested behavioral models are partitioned into nonstrategic behavioral models, all of which are elementary, and strategic behavioral models, none of which are elementary. Finally, in Section 7, we show that aggregating a finite number of elementary behavioral models results in a strictly more expressive

set of (not necessarily even minimally self-interested) models that still contains only nonstrategic behavioral models. We conclude in Section 8 with discussion and some future directions.

2. Background

We begin by briefly defining our formal framework and notation, discussing normal-form games, solution concepts, and behavioral models.

2.1. Normal-Form Games

In this work, we focus on unrepeated, simultaneous-move *normal-form games*. These games turn out to be perfectly general, in a mathematical sense: any game, including repeated games or dynamic “game trees”, can be represented as a normal-form game.

A normal-form game G is defined by a tuple (N, A, u) , where $N = \{1, \dots, n\}$ is a finite set of *agents*; $A = A_1 \times \dots \times A_n$ is the set of possible *action profiles*; A_i is the finite set of *actions* available to agent i ; and $u = \{u_i\}_{i \in N}$ is a set of *utility functions* $u_i : A \rightarrow \mathbb{R}$, each of which maps from an action profile to a utility for agent i . Agents may also randomize over their actions. It is standard in the literature to call such randomization a *mixed strategy*; however, for our purposes this terminology will be confusing, since it would lead us to discuss the strategies of nonstrategic agents. We thus instead adopt the somewhat nonstandard terminology *behavior* for this concept. We denote the set of agent i ’s possible behaviors by $S_i = \Delta^{|A_i|}$, and the set of possible *behavior profiles* by $S = S_1 \times \dots \times S_n$, where Δ^k denotes the standard k -simplex (the set $\left\{ \theta_0 u_0 + \dots + \theta_k u_k \mid \sum_{i=0}^k \theta_i = 1 \text{ and } \theta_i \geq 0 \text{ for all } i \right\}$), and hence $\Delta^{|X|}$ is the set of probability distributions over a discrete set X . Overloading notation, we represent the expected utility to agent i of a behavior profile $s \in S$ by $u_i(s)$. We use the notation s_{-i} to refer to the behavior profile of all agents except i , and (s_i, s_{-i}) to represent a full behavior profile.

2.2. Solution Concepts

A *solution concept* is a mapping from a game G to a behavior profile (or set of behavior profiles) that satisfies some criteria. Solution concepts can be interpreted descriptively, as a prediction of how agents will actually play a game. They can also be interpreted normatively, as a claim about how rational agents ought to play a game. (These two senses frequently overlap, as it is common to assume that agents will play a game rationally.) We will primarily be concerned with these solution concepts as formalizations of *strategic behavior* in games.

The foundational solution concept in game theory, and the most commonly used, is the Nash equilibrium.

Definition 1 (Nash equilibrium). Let $BR_i(s_{-i}) = \arg \max_{a_i \in A_i} u_i(a_i, s_{-i})$ denote the set of agent i ’s *best responses* to a behavior profile $s_{-i} \in S_{-i}$. A Nash equilibrium is a behavior profile in which every agent simultaneously best responds to all the other agents. Formally, $s^* \in S$ is a *Nash equilibrium* if

$$\forall i \in N, a_i \in A_i : s_i^*(a_i) > 0 \implies a_i \in BR_i(s_{-i}^*).$$

When agents play a Nash equilibrium, they must randomize independently. A correlated equilibrium relaxes this requirement, and allows for joint distributions of actions that are correlated.

Definition 2 (Correlated equilibrium). A correlated equilibrium is a distribution $\sigma \in \Delta^{|A|}$ over action profiles which satisfies the following for every agent $i \in N$ and every mapping $\xi_i : A_i \rightarrow A_i$:

$$\sum_{a \in A} \sigma(a) u_i(a_i, a_{-i}) \geq \sum_{a \in A} \sigma(a) u_i(\xi_i(a_i), a_{-i}).$$

Note that every Nash equilibrium s^* corresponds to a correlated equilibrium $\sigma^*(a) = \prod_{i \in N} s_i^*(a_i)$.

One important idea from behavioral economics is that people become more likely to make errors as the cost of making those errors decreases. This can be modeled by assuming that agents best respond *quantally*, rather than via strict maximization.

Definition 3 (Quantal best response). A (*logit*) *quantal best response* $QBR_i(s_{-i}; \lambda, G)$ by agent i to s_{-i} in game G is a behavior s_i such that

$$s_i(a_i) = \frac{\exp[\lambda \cdot u_i(a_i, s_{-i})]}{\sum_{a'_i} \exp[\lambda \cdot u_i(a'_i, s_{-i})]}, \quad (1)$$

where λ (the *precision* parameter) indicates how sensitive agents are to utility differences. When $\lambda = 0$, quantal best response is equivalent to uniform randomization; as $\lambda \rightarrow \infty$, quantal best response corresponds to best response in the sense that actions are played with positive probability only if they are best responses, i.e. $QBR_i(s_{-i}; \lambda)(a_i) > 0 \iff a_i \in BR_i(s_{-i})$.

The generalization from best response to quantal best response gives rise to a generalization of Nash equilibrium known as the *quantal response equilibrium* (“QRE”) [19].

Definition 4 (QRE). A *quantal response equilibrium* with precision λ is a behavior profile s^* in which every agent’s behavior is a quantal best response to the behaviors of the other agents; i.e., for all agents i ,

$$s_i^* = QBR_i^G(s_{-i}^*; \lambda).$$

2.3. Models of Agent Behavior

We now turn to what we term *behavioral models*, functions that return a behavior (a probability distribution over a single agent’s action space) for every given game. They differ from solution concepts because they represent a prediction of a *single agent’s* behavior, rather than a prediction about all agents’ joint behavior. Profiles of behavioral models can thus be seen as solution concepts that always encode a single product distribution over a given set of individual behaviors. We consider what can be said about the profiles of behavioral models induced by existing solution concepts in Section 4.

Before we can define behavioral models, we must introduce some basic notation. Let:

- \mathcal{G} denote the space of all normal-form games;
- $\mathbb{R}^* = \bigcup_{k=1}^{\infty} \mathbb{R}^k$ denote the space of all finite vectors;
- $(\mathbb{R}^*)^* = \bigcup_{k=1}^{\infty} (\mathbb{R}^*)^k$ denote the space of all finite-sized, finite-dimensional tensors;
- $\Delta^* = \bigcup_{k=1}^{\infty} \Delta^k$ denote the space of all finite standard simplices; and
- $A_{G,i}$ denote player i ’s action space in game G .

Definition 5 (Behavioral models). A *behavioral model* is a function $f_i : \mathcal{G} \rightarrow \Delta^*$; $f_i(G) \in \Delta^{|A_{G,i}|}$ for all games G . We use a function name with no agent subscript, such as f , to denote a *profile of behavioral models* with one function f_i for each agent. We write $f(G)$ to denote the behavior profile that results from applying each f_i to G .

Much work in behavioral game theory proposes behavioral models rather than solution concepts (though the formal definition of a behavioral model is our own). One key idea from that literature is that humans can only perform a limited number of steps of strategic reasoning, or equivalently that they only reason about higher-order beliefs up to some fixed, maximum order.

We begin with the so-called *level- k* model [21, 9]. Unlike Nash equilibrium, correlated equilibrium, and quantal response equilibrium, all of which describe fixed points, the level- k model is computed via a finite number of best response calculations. Each agent i is associated with a level $k_i \in \mathbb{N}$, corresponding to the number of steps of reasoning the agent is able to perform. A level-0 agent plays nonstrategically (i.e., without reasoning about its opponent); a level- k agent (for $k \geq 1$) best responds to the belief that all other agents are level- $(k-1)$. The level- k model implies a distribution over play for all agents when combined with a distribution $D \in \Delta^*$ over levels.

Definition 6 (Level- k prediction). Fix a distribution $D \in \Delta^*$ over levels and a level-0 behavior $s^0 \in S$. Then the *level- k behavior* for an agent i is defined as

$$s_i^k(a_i) = \begin{cases} 1/|BR_i(s_{-i}^{k-1})| & \text{if } a_i \in BR_i(s_{-i}^{k-1}), \\ 0 & \text{otherwise.} \end{cases}$$

The *level- k prediction* $\pi^{Lk} \in S$ for a game G is the average of the behavior of the level- k strategies weighted by the frequency of the levels,

$$\pi_i^{Lk}(a_i) = \sum_{k=0}^{\infty} D(k) s_i^k(a_i).$$

Cognitive hierarchy [6] is a very similar model in which agents respond to the distribution of lower-level agents, rather than believing that every agent performs exactly one step less of reasoning.

Definition 7 (Cognitive hierarchy prediction). Fix a distribution $D \in \Delta^*$ over levels and a level-0 behavior $s^0 \in S$. Then the *level- k hierarchical behavior* for an agent i is

$$\pi_i^k(a_i) = \begin{cases} 1/|BR_i(\pi_{-i}^{0:k-1})| & \text{if } a_i \in BR_i(\pi_{-i}^{0:k-1}), \\ 0 & \text{otherwise,} \end{cases}$$

where $\pi^0 = s^0$ and $\pi_i^{0:k-1}(a_i) = \sum_{m=0}^{k-1} D(m) \pi_i^m(a_i)$. The *cognitive hierarchy prediction* is again the average of the level- k hierarchical strategies weighted by the frequencies of the levels,

$$\pi_i^{CH}(a_i) = \sum_{k=0}^{\infty} D(k) \pi_i^k(a_i).$$

As we did with quantal response equilibrium, it is possible to generalize these iterative solution concepts by basing agents' behavior on quantal best responses rather than best responses. The resulting models are called *quantal level- k* and *quantal cognitive hierarchy* [e.g., 29, 33].

3. Strategic Behavioral Models

As discussed in the introduction, there is general qualitative agreement in the literature that strategic agents act to maximize their own utilities based on explicit probabilistic beliefs about the actions of other agents. Our ultimate goal is to characterize behavioral models that are unambiguously nonstrategic; thus, to strengthen our results, we adopt a somewhat more expansive notion of strategic behavior. Specifically, we define an agent as strategic if it satisfies two conditions, which we call (1) other responsiveness and (2) dominance responsiveness. These conditions require that the agent chooses actions both (1) with at least some dependence on others' payoffs; and (2) with at least some concern for its own payoffs.

Definition 8 (Strategic behavioral model). A behavioral model f_i is *strategic* if it is both other responsive and dominance responsive.

The key feature of strategic agents is that they take account of the incentives of other agents when choosing their own actions. To capture this intuition via the weakest possible necessary condition, we say that an agent is other responsive if its behavior is ever influenced by changes (only) to the utilities of other agents.

Definition 9 (Other responsiveness). A behavioral model f_i is *other responsive* if there exists any pair of games $G = (N, A, u)$ and $G' = (N, A, u')$ such that $u_i(a) = u'_i(a)$ for all $a \in A$, but $f_i(G) \neq f_i(G')$.

It is also traditional to assume that agents always act to maximize their expected utilities. This assumption is too strong for our purposes; for example, we want to allow for deviations from perfect utility maximization such as quantal best response. However, it does not seem reasonable to call an agent strategic if it pays no attention whatsoever to its own payoffs. We thus introduce a concept that we call *dominance responsiveness*, which is a considerably weaker sense in which an agent might show concern for its own payoffs. The condition requires only that gross changes in agents' own incentives will cause them to change their behavior. More formally, we say that a behavioral model is dominance responsive if for every pair of games $G, G' \in \mathcal{G}$ such that an action is strictly dominant in one game and strictly dominated in another, then the behavioral model does not behave identically in both games.

Definition 10 (Strict dominance). In a game $G = (N, A, u)$, an action $a_i \in A_i$ *strictly dominates* an action $a'_i \in A_i$ if and only if $u_i(a_i, a_{-i}) > u_i(a'_i, a_{-i}) \quad \forall a_{-i} \in A_{-i}$. If an action a_i^* strictly dominates every other action in game G , then it is *strictly dominant* in game G .

Definition 11 (Dominance responsiveness). A pair of games $G = (N, A, u)$ and $G' = (N, A, u')$ are *dominance reversed for agent $i \in N$* if there exist $a_i, a'_i \in A_i$ such that a_i is strictly dominant in G , but a'_i strictly dominates a_i in G' . A behavioral model $f_i : \mathcal{G} \rightarrow \Delta^*$ is *dominance responsive* if for every pair of games $G = (N, A, u)$ and $G' = (N, A, u')$ that are dominance reversed for i , $f_i(G) \neq f_i(G')$.

Observe that the two games G and G' are required to have nothing in common beyond their sizes and their use of the same names for the actions; the latter is to enable the statement $f_i(G) \neq f_i(G')$ to make sense. (Each evaluation of f_i is a vector; the inequality compares them element-wise, pairing up actions with the same names.) Also observe that we do not require i to play the dominant action in G or to do anything in particular in G' ; i is simply required to change behavior in some way when a strictly dominated action

changes to become dominated (and the rest of the game changes in an arbitrary way). Again, this is meant to be an extremely weak condition capturing the notion that an agent responds to its own payoffs.

4. Existing Solution Concepts are Strategic

We now demonstrate that our definition of strategic behavioral models does more than describe qualitative patterns of behavior that have been called “strategic” in the past: it also formally captures the predictions of various solution concepts both from classical game theory and from behavioral game theory (Nash equilibrium, correlated equilibrium, quantal response equilibrium, level- k , cognitive hierarchy, and quantal cognitive hierarchy).

We begin with a technical definition that is used in the proofs in this section: another notion of self-interest that we call self responsiveness.

Definition 12 (Self responsiveness). A behavioral model $f_i : \mathcal{G} \rightarrow \Delta^*$ is *self responsive* if for any game $G = (N, A, u)$, there exists a game $G' = (N, A, u')$ such that

1. $f_i(G') \neq f_i(G)$, and
2. $u'_j(a) = u_j(a)$ for all $a \in A$ and $j \neq i$.

We now show that quantal best response to a self responsive behavioral model is strategic.

Theorem 13. Fix a set $\{\lambda_i > 0\}_{i \in N}$ and a profile of behavioral models f that satisfies

$$f_i(G) = QBR_i(f_{-i}(G); \lambda_i, G)$$

for all games G and players i . Then every behavioral model f_i is strategic.

Proof. Suppose that f satisfies the above for some $\{\lambda_i > 0\}_{i \in N}$. We then prove the claim in two parts. In Part 1 we show that every behavioral model in f must be self responsive. In Part 2, we show that every behavioral model in f is strategic.

Part 1: For any profile of behavioral models f and $\lambda > 0$, the behavioral model $q_i(G) = QBR_i(f_{-i}(G); \lambda, G)$ is self responsive.

We prove this by showing how to construct a game G' satisfying Definition 12 from any game G . Fix an arbitrary $G = (N, A, u)$ and $i \in N$. One of agent i 's actions must be assigned weakly less probability by q_i than all others. Call this action $a_i^- \in \arg \min_{a_i \in A_i} f_i(G)(a_i)$. Now construct u' as follows:

$$u'_j(a) = \begin{cases} 1 & \text{if } j = i \text{ and } a_i = a_i^-, \\ 0 & \text{if } j = i \text{ and } a_i \neq a_i^-, \\ u_j(a) & \text{otherwise.} \end{cases}$$

Clearly G' differs from G only in i 's payoffs. Action a_i^- was assigned weakly less probability by $q_i(G)$ than any other action. However, a_i^- is strictly dominant in G' , and hence must be assigned strictly greater probability than any other action by $q_i(G')$. Therefore, $q_i(G) \neq q_i(G')$, and q_i is self responsive.

Part 2: For any $\lambda > 0$ and profile f of self responsive behavioral models, the behavioral model $q_i(G) = QBR_i(f_{-i}(G); \lambda, G)$ is strategic.

First, observe that strictly dominant actions always have higher expected utility than the actions that they dominate, and hence quantal best response to any behavioral model is dominance responsive, since higher expected utility actions are always played with higher probability.

It remains only to show that quantal response is other responsive. Consider the following game G .

	L	R
U	1, 1	0, 2
D	0, 2	1, 1
	G	

Let i be the row player and j be the column player. Choose a game G' with $f_j(G) \neq f_j(G')$ that differs only in j 's payoffs; such a game is guaranteed to exist because f_j is self responsive. Notice that $u_i(U, f_j(G)) = f_j(G)(L)$. Since $f_j(G) \neq f_j(G')$, $u_i(U, f_j(G)) \neq u_i(U, f_j(G'))$, and therefore $q_i(f_j(G))(U) \neq q_i(f_j(G'))(U)$. \square

Corollary 14. *All of QRE, Nash equilibrium, and correlated equilibrium are (profiles of) strategic behavioral models.*

Proof. Immediate from Theorem 13 and the observation that best response is a special case of quantal best response. \square

Theorem 15. *All of level- k , cognitive hierarchy, and quantal cognitive hierarchy are strategic behavioral models for agents of level 1 or level 2 (and higher).*

Proof. Let f^0 be a profile of behavioral models and $\lambda_i^1, \lambda_i^2 > 0$ for all $i \in N$. Choose behavioral model profiles f^1, f^2 satisfying $f_i^1(G) = QBR_i(f_{-i}^0(G); \lambda_i^1, G)$ and $f_i^2(G) = QBR_i(f_{-i}^1(G); \lambda_i^2, G)$ for all games G and players i .

If all of the behavioral models in f^0 are self responsive, then by the argument in Part 2 of the proof of Theorem 13, all of the behavioral models in f^1 are strategic. Thus, noting that best response is a special case of quantal best response, all of the listed iterative models are strategic for level 1.

Otherwise, all of the behavioral models in f^1 are self responsive by the argument in Part 1 of the proof of Theorem 13, and thus by the argument in Part 2 of the proof of Theorem 13, all of the behavioral models in f^2 are strategic. Thus, noting that best response is a special case of quantal best response, all of the listed iterative models are strategic for level 2. \square

5. Elementary Behavioral Models

Our main task in this paper is to separate nonstrategic behavior from strategic behavior. Now that we have formally defined the latter, we can introduce a class of behavioral models, called *elementary* models, that we will ultimately show are always nonstrategic. Observe that an agent reasoning strategically needs to account both for its own payoffs (in order to be dominance responsive) and for others' payoffs (in order to be other responsive); thus, it must evaluate each outcome in multiple terms. Our key idea is thus to require that nonstrategic behavior independently “scores” each outcome using a single number. In this section, we formalize such a notion and illustrate its generality via examples of how it can be used to encode previously proposed “nonstrategic” behaviors.

5.1. Defining Elementary Behavioral Models

The formal definition of elementary behavioral models is unfortunately more complex than the intuition we just gave. The reason is that any tuple of k real values can be encoded into a single real number; in information economics this is referred to as *dimension smuggling* [e.g., 22]. Without ruling out dimension smuggling, therefore, a restriction that nonstrategic agents rely on only a single number would lack any force. We thus define a *no smuggling* criterion; it depends in turn on the concept of a *dictatorial function*.

Definition 16 (Dictatorial function). A function $\varphi : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is *dictatorial* if its value is completely determined by a single input: $\exists i \in \{1, \dots, m\}$ such that $\forall x, x' \in \mathbb{R}^{n-1}, \forall a \in \mathbb{R}, \varphi(x_1, \dots, x_{i-1}, a, x_i, \dots, x_{n-1}) = \varphi(x'_1, \dots, x'_{i-1}, a, x'_i, \dots, x'_{n-1})$.

This class of functions takes its name from the social choice condition from which it is inspired; one input to the function φ acts as a dictator over φ 's output.

Our no smuggling condition says that either a function is dictatorial or that for every input dimension (i.e., action chosen by each player), there exist inputs for the function differing in that dimension that produce the same output.

Definition 17 (No smuggling). A function $\varphi : \mathbb{R}^* \rightarrow \mathbb{R}^n$ satisfies *no smuggling* iff either φ is dictatorial or for every $m > n$ and for every $i \in \{1, \dots, m\}$, there exist $x, x' \in \mathbb{R}^m$ such that $x_i \neq x'_i$ and $\varphi(x) = \varphi(x')$.

The no smuggling condition considers functions that reduce a vector of numbers to a smaller number of dimensions. For our purposes at the moment, it is helpful to imagine $n = 1$ (the vector is reduced to a single real number), though we will appeal to the general case in Section 7. The condition requires that φ summarizes its input in a meaningful sense, rather than simply encoding all of the original numbers into the infinite number of digits of a single real number. Observe that if φ is one to one (i.e., if every value in φ 's domain maps to a different value in its range) then it can be inverted, meaning that it performs dimension smuggling. The alternative is that there be at least one pair of inputs that produce the same output. For the sake of convenience, we impose a slightly stronger condition, that either there exists at least one such pair of vectors differing in each input dimension or the function is dictatorial. (That is, for the input dimension that completely determines a dictatorial function, we do not require the existence of a pair of inputs that produce the same output.)

Definition 18 (Elementary behavioral model). A behavioral model $f_i : \mathcal{G} \rightarrow \Delta^*$ is *elementary* if it can be represented as $h(\Phi(G))$, where

1. $\forall G = (N, A, u) \in \mathcal{G}, \forall a \in A, \Phi(G)_a = \varphi(u(a))$ for every action profile a ,
2. $\varphi : \mathbb{R}^* \rightarrow \mathbb{R}$ satisfies no smuggling, and
3. $h : (\mathbb{R}^*)^* \rightarrow \Delta^*$ is an arbitrary function; we use it to map from \mathbb{R}^A to $\Delta^{|A_i|}$.

For convenience, when condition 1 holds we refer to Φ as the *potential map* for φ .

That is, an elementary behavioral model works as follows. First, given an arbitrary game $G = (N, A, u)$, and for each action profile $a \in A$, we apply the same no-smuggling function φ to the $|N|$ -tuple of real values $\langle u_1(a), \dots, u_{|N|}(a) \rangle$, producing in each case a single real value. We represent all of these real values in a mapping we call Φ ; this *potential map* is a function of the same size as each of the utility functions. We then apply an arbitrary function h to $\Phi(G)$, producing a probability distribution over A_i .

5.2. Examples of Elementary Behavioral Models

To demonstrate the generality of elementary behavioral models, we show how to encode each of the candidate level-0 behavioral models that we proposed in our past work [32, 34]. (Thus, although that work only appealed to intuition, we can now conclude that these behavioral models are all nonstrategic.)

We begin with the simplest behavioral models: those that depend only on a given agent i 's utilities u_i .

Example 19 (Maxmax behavioral model). A maxmax action for agent i is an action giving rise to the best best case. An agent that wishes to maximize its possible payoff will play a maxmax action. The *maxmax behavioral model* $f_i^{\max\max}(G)$ uniformly randomizes over all of i 's maxmax actions in G .

$$f_i^{\max\max}(G)(a_i) \propto \begin{cases} 1 & \text{if } a_i \in \arg \max_{a'_i \in A_i} \max_{a_{-i} \in A_{-i}} u_i(a'_i, a_{-i}), \\ 0 & \text{otherwise.} \end{cases}$$

Example 20 (Maxmin behavioral model). A maxmin action for agent i is the action with the best worst-case guarantee. This is the safest action to play against hostile agents. The *maxmin behavioral model* $f_i^{\max\min}(G)$ uniformly randomizes over all of i 's maxmin actions in G :

$$f_i^{\max\min}(G)(a_i) \propto \begin{cases} 1 & \text{if } a_i \in \arg \max_{a'_i \in A_i} \min_{a_{-i} \in A_{-i}} u_i(a'_i, a_{-i}), \\ 0 & \text{otherwise.} \end{cases}$$

Example 21 (Minimax regret behavioral model). Following Savage [27], for each action profile, an agent has a possible *regret*: how much more utility could the agent have gained by playing the best response to the other agents' actions? Each of the agent's actions is therefore associated with a vector of possible regrets, one for each possible profile of the other agents' actions. A minimax regret action is an action whose maximum regret (in the vector of possible regrets) is minimal. The *minimax regret behavioral model* $f_i^{\text{mmr}}(G)$ uniformly randomizes over all of i 's minimax regret actions in G . That is, if

$$r(a_i, a_{-i}) = u_i(a_i, a_{-i}) - \max_{a_i^* \in A_i} u_i(a_i^*, a_{-i})$$

is the regret of agent i in action profile (a_i, a_{-i}) , then

$$f_i^{\text{mmr}}(G)(a_i) \propto \begin{cases} 1 & \text{if } a_i \in \arg \min_{a'_i \in A_i} \max_{a_{-i} \in A_{-i}} r(a'_i, a_{-i}), \\ 0 & \text{otherwise.} \end{cases}$$

Because each of the maxmax, maxmin, and minimax regret behavioral models depends only on agent i 's payoffs, we can set $\varphi(u) = u_i$ in each case; φ satisfies no smuggling because it is dictatorial. The encodings differ only in their choice of h . These vary in their complexity (e.g., maxmax: simply uniformly randomize over all actions that tie for corresponding to the largest potential value; for minimax regret it is necessary to compute a best response for each action profile). However, recall that h is an arbitrary function; thus, this is not a problem for our encoding.

Other behavioral models depend on both agents' utilities, and so require different φ functions.

Example 22 (Max welfare behavioral model). A *max welfare* action is part of some action profile that maximizes the sum of agents' utilities. The *max welfare behavioral model* $f_i^{\max\text{welfare}}(G)$ uniformly randomizes over max welfare actions in G :

$$f_i^{\max\text{welfare}}(G)(a_i) \propto \begin{cases} 1 & \text{if } a_i \in \arg \max_{a'_i \in A_i} \max_{a_{-i} \in A_{-i}} \sum_{j \in N} u_j(a'_i, a_{-i}), \\ 0 & \text{otherwise.} \end{cases}$$

We can encode the efficient behavioral model as elementary by setting $\varphi(u) = \sum_j u_j$; φ satisfies no smuggling because it is continuous. We then define h to uniformly randomize over all actions that tie for corresponding to the largest potential value.

Example 23 (Fair behavioral model). Let the unfairness of an action profile be the difference between the maximum and minimum payoffs among the agents under that action profile:

$$d(a) = \max_{i, j \in N} u_i(a) - u_j(a).$$

Then a “fair” outcome minimizes this difference in utilities. A *fair* action is part of a minimally unfair action profile. The *fair behavioral model* $f_i^{\text{fair}}(G)$ uniformly randomizes over fair actions:

$$f_i^{\text{fair}}(G)(a_i) \propto \begin{cases} 1 & \text{if } a_i \in \arg \min_{a'_i \in A_i} \min_{a_{-i} \in A_{-i}} d(a'_i, a_{-i}), \\ 0 & \text{otherwise.} \end{cases}$$

We can encode the fair behavioral model as elementary by setting $\varphi(u) = \max_{j,k}(u_j - u_k)$; again, φ satisfies no smuggling because it is continuous. We then define h to uniformly randomize over all actions that tie for corresponding to the smallest potential value.

Finally, we note that all of the examples just given are binary: actions are either fair/maxmin/etc or they are not. By changing only h , we could similarly construct continuous variants of each concept in which actions that achieve nearly maximal potentials are played nearly as often by the behavioral model, etc.

6. Elementary Behavioral Models are Nonstrategic

We are now ready to show that elementary behavioral models are never strategic. This result is important because it achieves our key goal of distinguishing strategic from nonstrategic behavioral models via a formal mathematical criterion, rather than relying on the intuitive sense that a model “depends on” an explicit model of an opponent’s behavior. In fact, we do a bit better than simply showing that elementary models are nonstrategic: we show that the space of dominance responsive behavioral models is exactly partitioned into elementary models and strategic models.

Theorem 24. *A dominance responsive behavioral model is strategic if and only if it is not elementary.*

Proof. If direction: no elementary behavioral model is strategic. Suppose for contradiction that elementary behavioral model $f_i(G) = h_i(\Phi(G))$ is strategic, where Φ is the potential map for φ . There are two ways in which φ can satisfy the no-smuggling condition:

1. *i is a dictator for φ .* Because f_i is other responsive, there exist $G = (N, A, u)$ and $G' = (N, A, u')$ with $u_i(a) = u'_i(a)$ for all $a \in A$, such that $f_i(G) \neq f_i(G')$. But since i is a dictator for φ , $\Phi(G) = \Phi(G')$, and hence $f_i(G) = f_i(G')$, a contradiction.
2. *For all $m > 1$, there exist $x, x' \in \mathbb{R}^m$ such that $x_i \neq x'_i$ and $\varphi(x) = \varphi(x')$.* Let $x, x' \in \mathbb{R}^2$ be two such values; we use them to construct 2-player games G_1 and G_2 in which i has two actions $\{U, D\}$ and the other agent has actions $\{L, R\}$.

$$\begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\ \begin{array}{c} U \\ D \end{array} & \begin{array}{|cc|} \hline x' & x' \\ \hline x & x \\ \hline \end{array} \end{array} \quad \begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\ \begin{array}{c} U \\ D \end{array} & \begin{array}{|cc|} \hline x & x \\ \hline x' & x' \\ \hline \end{array} \end{array}$$

$G_1 \qquad G_2$

Note that U is a strictly dominant action for i in G_1 and a strictly dominated action for i in G_2 . By dominance responsiveness, $f_i(G_1) \neq f_i(G_2)$. But $\varphi(x) = \varphi(x')$, and since x and x' are the only payoff tuples that occur in either G_1 or G_2 , $\Phi(G_1) = \Phi(G_2)$ and hence $f_i(G_1) = f_i(G_2)$, a contradiction.

Only-if direction: if a dominance responsive behavioral model f_i is not strategic, then it is elementary. We show how to represent f_i by constructing appropriate φ and Φ functions. Since f_i is dominance responsive but not strategic, it is not other responsive. Therefore, for every pair of games $G = (N, A, u)$ and $G' = (N, A, u')$ with $u_i(a) = u'_i(a)$ for all $a \in A$, $f_i(G) = f_i(G')$. Define $\varphi(x) = x_i$, which clearly satisfies no smuggling. Let $h_i(\Phi(G)) = f_i(z(\Phi(G)))$, where $z : \mathbb{R}^A \rightarrow \mathcal{G}$ is a function that returns a game with the utilities of i set to its argument and the utilities of the other players set to 0. Since differences in the other agents’ utilities never change the output of f_i , $h_i(\Phi(G)) = f_i(z(\Phi(G))) = f_i(G)$ for all G . \square

7. Finite Aggregations of Elementary Models

We now consider behavioral models that are constructed by drawing together the predictions of multiple elementary models. For example, we might build a behavioral model from some convex combination of the predictions of the elementary behavioral models defined in Section 5.2 (as, indeed, we did in our past work [32, 34]). Our key result in this section is that the class of such models is strictly larger than the class of elementary behavioral models. However, this larger class still consists entirely of nonstrategic models, under an appropriate strengthening of our no smuggling condition.

How can this be—didn't we already characterize all nonstrategic behavioral models? Not quite: recall that Theorem 24 partitioned the set of *dominance responsive* behavioral models into strategic and elementary behavioral models. In addition, some elementary models are non-dominance-responsive. The set of non-dominance-responsive behavioral models also includes the extra models made possible by taking finite aggregations. This section thus introduces a natural class of nonstrategic behavior that sometimes violates our minimal notion of self interest without allowing completely arbitrary behavior.

Definition 25 (Finite aggregations of elementary behavioral models). A behavioral model g_i is a *finite aggregation of elementary behavioral models* if it can be represented as

$$g_i(G) = h(f_i^1(G), \dots, f_i^K(G)),$$

where h is an arbitrary function, $K \in \mathbb{N}$, and for all $1 \leq k \leq K$, f_i^k is an elementary behavioral model.

The class of finite aggregations of elementary behavioral models is strictly larger than the class of elementary behavioral models.

Theorem 26. *Finite aggregations of elementary behavioral models are not necessarily elementary themselves.*

Proof. It suffices to exhibit two elementary behavioral models whose linear combination is not elementary. Let $f_i^{\max\max}$ be the maxmax behavioral model defined in Example 19; let $f_i^{\max\text{welfare}}$ be the max welfare behavioral model defined in Example 22. Let $\alpha \in (0, 1)$ and let $g_i(G) = h(f_i^{\max\max}(G), f_i^{\max\text{welfare}}(G)) = \alpha f_i^{\max\max}(G) + (1 - \alpha) f_i^{\max\text{welfare}}(G)$.

Suppose for contradiction that g_i is elementary. Since g_i is a function of all agents' utilities, its potential φ is not dictatorial for any agent. Therefore, from no smuggling, there must exist two utility vectors x, x' satisfying $x'_i > x_i$ and $\varphi(x') = \varphi(x)$. Let $\delta = (x'_i - x_i)/2$, $y = \langle x_i + \delta/2, \max\{x_j, x'_j\} + 2\delta \rangle$, and z be any utility tuple satisfying $z_i < x_i$ and $z_j < \min\{x_j, x'_j\}$. Now consider the following two-player games G_3 and G_4 , in which i is the row player.

$$\begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\ \begin{array}{c} U \\ D \end{array} & \begin{array}{|c|c|} \hline x & z \\ \hline z & y \\ \hline \end{array} \end{array} \quad \begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\ \begin{array}{c} U \\ D \end{array} & \begin{array}{|c|c|} \hline x' & z \\ \hline z & y \\ \hline \end{array} \end{array}$$

$G_3 \qquad G_4$

Since the two games differ only in outcome $\langle U, L \rangle$, and $\varphi(u(U, L)) = \varphi(u'(U, L))$, it must be that $g_i(G_3) = g_i(G_4)$. But $g_i(G_3)(U) = 0 \neq \alpha = g_i(G_4)(U)$, yielding our contradiction. \square

We need to strengthen our no-smuggling condition for behavioral models based on finite aggregations of elementary models, because the existence of multiple potential functions and the fact that we allow arbitrary aggregation functions makes our previous definition too weak to prevent dimension smuggling.

Definition 27 (Joint no smuggling). A set of functions $\varphi_1, \dots, \varphi_K : \mathbb{R}^m \rightarrow \mathbb{R}^n$ for $m > Kn$ is *jointly no-smuggling* if the function $\psi(x) = [\varphi_1(x), \dots, \varphi_K(x)]$ satisfies the no-smuggling condition.

The joint no-smuggling condition is a benign assumption. In particular, it is satisfied by every set of continuous potentials, including all linear combinations of utilities (note that a behavioral model with a continuous potential need not be continuous itself).

This new definition suffices to give us our final result: that finite aggregations of elementary behavioral models are always nonstrategic.

Theorem 28. *Let F be a set of elementary behavioral models. If the potentials $\varphi_1, \dots, \varphi_{|F|}$ associated with the models in F are jointly no-smuggling, then no finite aggregation of F is strategic.*

Proof. The argument in the proof of the if direction of Theorem 24 can be applied to finite aggregations of elementary behavioral models in exactly the same way, since the joint no-smuggling condition guarantees the existence of utility vectors x, x' that are indistinguishable by the joint potential ψ but that have $x'_i > x_i$. \square

8. Discussion and Future Work

In this work, we proposed elementary behavioral models and their finite aggregations as mathematical characterizations of classes of nonstrategic decision rules. These classes are constructively defined, in the sense that membership of a rule is verified by demonstrating how to represent the rule in a specific form—as a function of the output of a non-smuggling potential map—rather than by proving that it cannot be represented as a response to probabilistic beliefs.

It is interesting to note that various special cases of strategic solution concepts are nonstrategic under our definition. For example, the equilibrium of a two-player zero-sum game can be computed by considering only the utility of a single agent, and hence the behavior for an equilibrium-playing player in such a game can be computed by an elementary behavioral model. Similarly, an equilibrium for a potential game can of course be computed in terms of outcome values computed by a potential function [20]. In repeated settings, many no-regret learning rules (which are guaranteed to converge to a coarse correlated equilibrium) can be executed by agents that only take account of their own utilities. One thing that these exceptions all have in common is that they are also *computationally* easy, unlike general ϵ -equilibrium, which is known to be hard in a precise computational sense [11, 8]. The equilibrium of a zero-sum game can be solved in polynomial time by a linear program; the equilibrium of a potential game can be found simply by finding the maximum of the potential function over all the outcomes. No-regret algorithms are cheap to run, requiring in each time period work that is linear in the number of actions, and converge rapidly to coarse correlated equilibrium. However, the connection between ease of computation and strategic simplicity is not an equivalence. For example, correlated equilibrium in general games can be computed in polynomial time by a linear program, but cannot in general be computed by an elementary behavioral model (see Corollary 14). An attractive future direction is to shed further light on the connection between computational and strategic simplicity.

We also observe that our characterization of nonstrategic behavior in this paper is a binary distinction: in the view we have advanced, a behavioral model is either nonstrategic or it is not. An intriguing question

for future work is whether such a distinction can be made more quantitative: i.e., is there a sense in which agents are nonstrategic to a greater or lesser degree that is distinct from the number of steps of strategic reasoning that they perform?

Acknowledgements

This work was funded in part by an NSERC E.W.R. Steacie Fellowship and an NSERC Discovery Grant. Part of this work was done at Microsoft Research New York while the first author was a postdoctoral researcher and the second author was a visiting researcher.

References

- [1] Adlakha, S. and Johari, R. (2013). Mean field equilibrium in dynamic games with strategic complementarities. *Operations Research*, 61(4):971–989.
- [2] Airiau, S. and Sen, S. (2003). Strategic bidding for multiple units in simultaneous and sequential auctions. *Group Decision and Negotiation*, 12(5):397–413.
- [3] Babaioff, M., Nisan, N., and Pavlov, E. (2004). Mechanisms for a spatially distributed market. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 9–20.
- [4] Bar-Isaac, H. and Ganuza, J.-J. (2005). Teaching to the top and searching for superstars. Technical report, New York University, Leonard N. Stern School of Business, Department of Economics.
- [5] Bernheim, B. (1984). Rationalizable Strategic Behavior. *Econometrica*, 52(4):1007–1028.
- [6] Camerer, C., Ho, T., and Chong, J. (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119(3):861–898.
- [7] Camerer, C. F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.
- [8] Chen, X. and Deng, X. (2006). Settling the complexity of two-player nash equilibrium. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*, pages 261–272.
- [9] Costa-Gomes, M., Crawford, V., and Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5):1193–1235.
- [10] Crawford, V. P., Costa-Gomes, M. A., Iriberri, N., et al. (2010). Strategic thinking. Working paper.
- [11] Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. (2009). The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259.

- [12] Gerding, E. H., Robu, V., Stein, S., Parkes, D. C., Rogers, A., and Jennings, N. R. (2011). Online mechanism design for electric vehicle charging. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 811–818.
- [13] Ghosh, A. and Hummel, P. (2012). Implementing optimal outcomes in social computing: a game-theoretic approach. In *Proceedings of the 21st International Conference on World Wide Web*, pages 539–548.
- [14] Gilboa, I. and Samet, D. (1989). Bounded versus unbounded rationality: The tyranny of the weak. *Games and Economic Behavior*, 1(3):213–221.
- [15] Grabisch, M., Mandel, A., Rusinowska, A., and Tanimura, E. (2017). Strategic influence in social networks. *Mathematics of Operations Research*, 43(1):29–50.
- [16] Hartford, J. S., Wright, J. R., and Leyton-Brown, K. (2016). Deep learning for predicting human strategic behavior. In *Advances in Neural Information Processing Systems*, pages 2424–2432.
- [17] Lee, H. (2014). Algorithmic and game-theoretic approaches to group scheduling. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1709–1710.
- [18] Li, C. and Tesauro, G. (2003). A strategic decision model for multi-attribute bilateral negotiation with alternating. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 208–209.
- [19] McKelvey, R. and Palfrey, T. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38.
- [20] Monderer, D. and Shapley, L. S. (1996). Potential games. *Games and Economic Behavior*, 14(1):124–143.
- [21] Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5):1313–1326.
- [22] Nisan, N. and Segal, I. (2006). The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224.
- [23] Pearce, D. (1984). Rationalizable Strategic Behavior and the Problem of Perfection. *Econometrica*, 52(4):1029–1050.
- [24] Roth, A. E. and Ockenfels, A. (2002). Last-minute bidding and the rules for ending second-price auctions: Evidence from ebay and amazon auctions on the internet. *American economic review*, 92(4):1093–1103.
- [25] Rubinstein, A. (1998). *Modeling bounded rationality*. MIT press.
- [26] Sandholm, T. W. and Lesser, V. R. (2001). Leveled commitment contracts and strategic breach. *Games and Economic Behavior*, 35(1-2):212–270.

- [27] Savage, L. (1951). The Theory of Statistical Decision. *Journal of the American Statistical Association*, 46(253):55–67.
- [28] Shoham, Y. and Leyton-Brown, K. (2008). *Multiagent Systems: Algorithmic, Game-theoretic, and Logical Foundations*. Cambridge University Press.
- [29] Stahl, D. and Wilson, P. (1994). Experimental evidence on players’ models of other players. *Journal of Economic Behavior and Organization*, 25(3):309–327.
- [30] Train, K. (2009). *Discrete Choice Methods with Simulation*. Cambridge University Press.
- [31] Walsh, W. E., Wellman, M. P., and Ygge, F. (2000). Combinatorial auctions for supply chain formation. In *Proceedings of the 2nd ACM conference on Electronic commerce*, pages 260–269.
- [32] Wright, J. R. and Leyton-Brown, K. (2014). Level-0 meta-models for predicting human behavior in games. In *Proceedings of the ACM Conference on Economics and Computation (EC’14)*, pages 857–874.
- [33] Wright, J. R. and Leyton-Brown, K. (2017). Predicting human behavior in unrepeated, simultaneous-move games. *Games and Economic Behavior*, 106:16–37.
- [34] Wright, J. R. and Leyton-Brown, K. (2019). Level-0 models for predicting human behavior in games. *Journal of Artificial Intelligence Research*, 64:357–383.