

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Kline, Brendan

Article An empirical model of non-equilibrium behavior in games

Quantitative Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Kline, Brendan (2018) : An empirical model of non-equilibrium behavior in games, Quantitative Economics, ISSN 1759-7331, The Econometric Society, New Haven, CT, Vol. 9, Iss. 1, pp. 141-181, https://doi.org/10.3982/QE647

This Version is available at: https://hdl.handle.net/10419/195565

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.



https://creativecommons.org/licenses/by-nc/4.0/

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.





An empirical model of non-equilibrium behavior in games

BRENDAN KLINE

Department of Economics, University of Texas at Austin

This paper studies the identification and estimation of the decision rules that individuals use to determine their actions in games, based on a structural econometric model of non-equilibrium behavior in games. The model is based primarily on various notions of limited strategic reasoning, allowing multiple modes of strategic reasoning and heterogeneity in strategic reasoning across individuals and within individuals. The paper proposes the model and provides sufficient conditions for point identification of the model. Then the model is estimated on data from an experiment involving two-player guessing games. The application illustrates the empirical relevance of the main features of the model.

KEYWORDS. Games, heterogeneity, identification, non-equilibrium, strategic reasoning.

JEL CLASSIFICATION. C1, C57, C72.

1. INTRODUCTION

In game theory, different solution concepts and decision rules make different predictions about how players determine their actions given their utility functions. The Nash equilibrium solution concept (e.g., Nash (1950)) is the most common prediction about how players behave, but theory also provides other solution concepts and decision rules that make different predictions about how players behave. Indeed, there is considerable empirical evidence of behavior that does not conform to the predictions of Nash equilibrium (e.g., Camerer (2003)).

Because of the central role of game theory in economics and other disciplines, it is important to conduct empirical investigations that evaluate these solution concepts and decision rules. The credibility of predictions based on game theory models depends on the credibility of the solution concept or decision rule that generates the predictions because, by definition, different solution concepts and decision rules can generate different predictions even for the same specification of the utility functions. Further, the lit-

Brendan Kline: brendan.kline@austin.utexas.edu

Thanks to many individuals and seminar audiences at Harvard University/MIT, London School of Economics, Northwestern University, University of Texas at Austin, Vanderbilt University, the Cowles Foundation Conference on Econometrics, Texas Econometrics Camp 2015, Texas Economic Theory Camp 2015, and the 2015 World Congress of the Econometric Society for helpful comments and suggestions. Thanks to the co-editor and two anonymous referees for further helpful comments and suggestions. Finally, thanks to Costa-Gomes and Crawford (2006) for making their data publicly available. Any errors are mine.

Copyright © 2018 The Author. Quantitative Economics. The Econometric Society. Licensed under the Creative Commons Attribution-NonCommercial License 4.0. Available at http://www.qeconomics.org. https://doi.org/10.3982/QE647

erature¹ on estimation of the utility functions in game theory models involves assumptions on the solution concept. If the assumptions on the solution concept are false, then the resulting model is misspecified, raising concerns about the credibility of the empirical results.

With that broad motivation, this paper is concerned with a particular structural econometric model of non-equilibrium behavior in games. Rather than estimate the utility functions under the assumption that the econometrician knows the solution concept, as in the prior literature on the econometrics of games, this paper is concerned with estimating the solution concept(s) or decision rule(s) under the assumption that the econometrician knows the utility functions, as in the payoffs presented to subjects in an experiment. Similar empirical questions have been a major focus in the literature on experimental economics. The focus in this paper is on understanding the identification of the model. Often the models and empirical strategies used in experimental economics are point identified by relatively short arguments such that the identification problem is not the main focus of the papers. The model proposed in this paper leads to a more challenging identification problem. This paper proposes the model, establishes sufficient conditions for point identification of the model, and estimates the model on real data. Following the economic theory and experimental economics literatures, the model is based on two main classes of alternatives to Nash equilibrium relating to limited strategic reasoning.

Unanchored strategic reasoning is a model of limited strategic reasoning that can be viewed as an empirical implementation of ideas from the economic theory literature relating to rationalizability (e.g., Bernheim (1984) and Pearce (1984)) and iterated deletion of dominated strategies.² The model includes different numbers of steps of unanchored strategic reasoning, interpreted as different levels of sophistication of strategic reasoning. In most games, a set of actions are consistent with any given number of steps of unanchored strategic reasoning and a given action can be consistent with multiple different numbers of steps of unanchored strategic reasoning. Therefore, it is not possible to infer the number of steps of unanchored strategic reasoning that an individual uses by inspecting whether the action taken by that individual is equal to that predicted by a particular number of steps of unanchored strategic reasoning. For example, the observation that an individual uses a particular action could be consistent with that individual using

¹Papers in that literature (among papers that study complete information), typically based on Nash equilibrium, include Tamer (2003), Aradillas-Lopez and Tamer (2008), Bajari, Hong, and Ryan (2010), Kline and Tamer (2012), Aradillas-Lopez and Rosen (2013), Dunker, Hoderlein, and Kaido (2013), Kline (2015, 2016), and Fox and Lazzati (2017). See de Paula (2013) for a review. Important identification results in experimental economics include Haile, Hortaçsu, and Kosenok (2008) (quantal response equilibrium model), and Gillen (2010) and An (2017) (level-*k* model in auctions).

²From the economic theory literature, rationalizability is equivalent to common knowledge of rationality and independence of actions across players (e.g., Tan and da Costa Werlang (1988)). In two-player games, rationalizability is also equivalent to infinitely many steps of iterated deletion of dominated strategies. See, for example, Tan and da Costa Werlang (1988) or Fudenberg and Tirole (1991). In any given game, rationalizability might be equivalent to a certain finite number of steps of iterated deletion of dominated strategies if additional strategies are no longer deleted in further iterations. But, in general, rationalizability requires infinitely many (or, at least, unbounded) steps of iterated deletion.

either zero or one step of unanchored strategic reasoning. This results in one part of the identification problem studied in this paper. The identification result establishes how it is possible to identify/estimate the number of steps of unanchored strategic reasoning that individuals use. Moreover, evidence of unanchored strategic reasoning is found in the empirical application.

The experimental economics literature also has models of limited strategic reasoning. Anchored strategic reasoning is a model of limited strategic reasoning otherwise known as the level-k model of thinking that is commonly used in the experimental game theory literature.³ In the level-k model of thinking, individuals who use zero steps of reasoning are "anchored" to a particular distribution of actions, usually the uniform distribution over the action space. Hence, this paper uses the term anchored strategic reasoning to refer to this decision rule. Then, somewhat similarly to unanchored strategic reasoning, individuals who use more than zero steps of anchored strategic reasoning best respond to the strategy used by an individual of the immediately lower number of steps of anchored strategic reasoning.

Rather than suppose that a single decision rule is responsible for generating all actions of all individuals, the model allows both across-individual and within-individual heterogeneity in the decision rule(s). Consequently, the goal of the model is to estimate how often individuals use each of the decision rules. In particular, the goal of the model is to estimate how often individuals use each number of steps of unanchored and/or anchored strategic reasoning.

Across-individual heterogeneity allows that different individuals use different decision rules, an important stylized fact from the experimental game theory literature. Similarly, within-individual heterogeneity allows that even a given individual uses multiple different decision rules, a contribution of the model in this paper. Prior empirical work in the related experimental game theory literature has been based on the assumption that each individual uses just one decision rule. In particular, the prior literature based on the level-k model of thinking generally characterizes individuals as a "level-1" thinker or a "level-2" thinker, and so on. The model in this paper allows that a given individual is characterized by the use of multiple decision rules, rather than just one decision rule, just as the overall population of individuals is characterized by the use of multiple decision rules, rather than just one decision rule. As discussed in Section 2.4, among other interpretations, within-individual heterogeneity can be given an interpretation similar to random utility models in single-agent decision problems, in the sense that the behavior of individuals in games may be described as arising from randomly selecting from a set of decision rules. Across-individual heterogeneity and within-individual heterogeneity have similar observable implications, since both involve the use of multiple decision rules. Therefore, heterogeneity in the decision rules results in another part of the identification problem studied in this paper. For example, suppose that there are decision rules A and B in the model. Based on data from individuals playing any given game,

³See, for example, Camerer (2003) or Crawford, Costa-Gomes, and Iriberri (2012) for a discussion of the related experimental literature, which includes, in particular (not exhaustive), Stahl and Wilson (1994, 1995), Nagel (1995), Ho, Camerer, and Weigelt (1998), Costa-Gomes, Crawford, and Broseta (2001), Camerer, Ho, and Chong (2004), Costa-Gomes and Crawford (2006), and Crawford and Iriberri (2007a, 2007b).

there can be observational equivalence between two distinct specifications of heterogeneity: (a) one type of individual always using A and another type of individual always using B, as in across-individual heterogeneity, and (b) each individual using both A and B, as in within-individual heterogeneity. As detailed by the arguments in Section 3.2, both specifications are such that some actions used in the data are consistent with decision rule A and other actions used in the data are consistent with decision rule B. The identification result establishes how to identify/estimate the heterogeneity in the use of multiple decision rules. The empirical application shows evidence of both acrossindividual and within-individual heterogeneity.

The paper establishes sufficient conditions for point identification of the unknown parameters. In the absence of such sufficient conditions, the paper shows by example that it can easily happen that the unknown parameters are not point identified. Many of the main sufficient conditions concern the structure of the games that experimental subjects are observed to play. Consequently, the identification results can guide experimental design. The range of experimental designs that have been used within experimental game theory are discussed, for example, in Camerer (2003) and Crawford, Costa-Gomes, and Iriberri (2012). One of the main sufficient conditions is that the econometrician observes each subject play multiple games. The identification result characterizes how many games the subjects must play as a function of the degree of across-individual heterogeneity.

Then the model is estimated using data that come from the two-person guessing game experiment in Costa-Gomes and Crawford (2006) to establish the empirical relevance of the results in the context of a well known and representative experimental design. The results suggest that both across-individual and within-individual heterogeneity, and unanchored strategic reasoning are important. For example, the most common type of subject in the experiment is estimated to comprise approximately 44% of the population: it uses zero steps of unanchored strategic reasoning with probability approximately 49% and one step of unanchored strategic reasoning with probability approximately 31%. It also uses anchored strategic reasoning and Nash equilibrium with cumulative probability approximately 21%. The identification results establish how it is possible to recover these parameters from the data. In contrast, related models in experimental game theory do not include unanchored strategic reasoning and are restricted to types of subjects who exclusively use one decision rule. Such models would, therefore, not capture all of the characteristics of the subjects.

In addition to the differences due to focusing on identification of the model in general rather than empirical results from a particular experiment, the model in this paper differs from prior models in experimental game theory. Those differences are the reason for the more challenging identification problem in this paper. In particular, allowing unanchored strategic reasoning and within-individual heterogeneity substantially complicates the identification problem, and the application shows that those features of the model are empirically relevant. As discussed further in Section 3, these two features of the model independently complicate the identification problem. Identifying the model that allows unanchored strategic reasoning is complicated even when restricting to a model without within-individual heterogeneity, and identifying the model that allows within-individual heterogeneity is complicated even when restricting to a model without unanchored strategic reasoning. Therefore, the identification result is a relevant contribution even if some but not all of those features are present in a particular application.

The rest of the paper is organized as follows. Section 2 sets up the model. Section 3 sets up the identification problem, and Section 4 establishes sufficient conditions for point identification. Section 5 reports the empirical application. Section 6 concludes. The Appendices, available in a supplementary file on the journal website, http:// qeconomics.org/supp/647/supplement.pdf, collect supplemental results, including derivation of the model likelihood (Appendix A), point identification of all model parameters except for the magnitude of computational mistakes (Appendix B), discussion of identification of the selection rule on unanchored strategic reasoning (Appendix C), the proofs of the point identification assumptions hold in the empirical application (Appendix E), and additional empirical results (Appendices F and G). Replication files are available in a supplementary file on the journal website, http://geconomics.org/supp/647/code_and_data.zip.

2. Model

2.1 Notation for the games

The goal of the model is to study strategic behavior in complete information games with continuous action spaces.⁴ The setup for game g is as follows.

(i) There are M_g players, indexed by $j = 1, 2, ..., M_g$. Note that "the player indexed by j" or just "player j" corresponds to the indexing of players in the game, and is not the same as subject j in the data set. Therefore, player j might alternatively be called, for example, the row player in the game.

(ii) The action of player *j* is a_j . The action space for player *j* in game *g* is the interval of real numbers $[\alpha_{Lg}(j), \alpha_{Ug}(j)]$ and, consequently, there is a continuous action space.

(iii) The utility function of player *j* in game *g* is $u_{jg}(a_1, \ldots, a_{M_g})$.

(iv) All of these facts are common knowledge among the players, so the game is complete information. Also, all of these facts are known by the econometrician.

As formalized in Section 2.7, the econometrician has data on the behavior of subjects in these games. There is an important distinction between player and subject. The term *subject* refers to an actual individual (e.g., an "experimental subject") in the real world. The term *player* refers to the more generic game theory concept. For example, a player could refer to the "row player" in a particular game. Consequently, when the experiment

⁴It is possible to specify a similar model for games with discrete action spaces and to identify such a model using an adaptation of the identification strategy for the model with continuous action spaces. Games with continuous action spaces provide more scope for different decision rules to make different predictions about the action an individual takes, which is necessary for identification.

has subjects play the games, subjects are assigned the roles of particular players in the games.

2.2 Decision rules

The model is concerned with recovering information about the solution concepts and decision rules that subjects use, based on observing the behavior of those subjects. By *solution concept*, this paper means a possibly set-valued mapping between the specification of a game and the set of strategies for all of the players. By *decision rule*, this paper means a possibly set-valued mapping between the specification of a game and the set of strategies for all of the players. By *decision rule*, this paper means a possibly set-valued mapping between the specification of a game and the set of strategies for an individual player. Each solution concept and decision rule can be viewed as making a set-valued prediction about behavior.⁵ In particular, following the literature on experimental game theory, this paper focuses on non-equilibrium solution concepts and decision rules. Even equilibrium solution concepts like Nash equilibrium can be viewed as making non-equilibrium predictions, in the sense of making predictions for each individual player. Consequently, a player can be said to use its part of a solution concept (e.g., Nash equilibrium) or can be said to use a certain decision rule, without consideration of the actual behavior of the other players in the game.

Sections 2.2.1-2.2.3 describe the decision rules included in the model. These decision rules are demonstrated by example in the empirical application in Section 5. An extended model that admits the possibility that subjects use an enlarged class of candidate decision rules could be identified by extending the identification strategy. However, it is not realistic to expect identification of a model that does not involve some sort of ex ante restriction on the class of candidate decision rules. An unrestricted model of nonequilibrium behavior is fundamentally unidentified. With no restrictions on the class of candidate decision rules, there are infinitely many decision rules that coincide with any observed behavior of any particular subject in the data but differ in their predictions about other games. Such decision rules are trivial to characterize: simply specify that they are mappings from the space of games to the space of strategies that specifically coincide with the games and corresponding actions observed in the data for that particular subject, but differ on different games not observed in the data. Since decision rules are mappings from the space of games to the space of strategies, estimating the decision rule without ex ante known restrictions on the space of decision rules would require the econometrician to observe the subjects' behavior in all games, which is effectively impossible.6

⁵Similar definitions have been used in the game theory literature; see, for example, Myerson (1991, pp. 88 and 107), Osborne and Rubinstein (1994, p. 2), or Aumann (2000, p. 57).

⁶The restrictions on the class of candidate decision rules provided by economic theory play a somewhat similar role in the identification strategy as do function space assumptions (e.g., finite-dimensional parametrization, continuity, etc.) in regression function estimation, in the sense that they make it possible to identify/estimate a high-dimensional object (i.e., either a decision rule, or function) based on somehow (indirectly) "observing" that object at a strict subset of its domain (i.e., a decision rule at a subset of games, or a function at a subset of its domain).

2.2.1 *Nash equilibrium* The Nash equilibrium solution concept predicts that players use strategies that are mutually best responses. According to Nash equilibrium, player *j* in game *g* uses a strategy σ_{jg} , with the property that σ_{jg} is a distribution supported on the set of solutions to

$$\max_{a_j\in[\alpha_{Lg}(j),\alpha_{Ug}(j)]}E_{\sigma_{-j,g}}(u_{jg}(a_1,\ldots,a_{M_g})),$$

where $[\alpha_{Lg}(j), \alpha_{Ug}(j)]$ is the action space for player *j* in game *g*, and the expectation notation indicates that a_{-j} are distributed according to the Nash equilibrium strategies of the other players in game *g* (i.e., according to $\sigma_{-j,g}$). The model is based on the assumption that there is a unique pure strategy Nash equilibrium that predicts that player *j* in game *g* takes action c_{jg} (NE), as is the typical case for games studied in the related experimental game theory literature.⁷ The notation for Nash is NE.

2.2.2 Unanchored strategic reasoning Unanchored strategic reasoning is a class of decision rules that are iteratively defined steps of increasingly sophisticated strategic reasoning related to iterated deletion of dominated strategies, particularly in two-player games,⁸ and rationalizability (e.g., Bernheim (1984) and Pearce (1984)). One contribution of this paper is to study the empirical relevance of unanchored strategic reasoning by providing a model in which it is possible to identify/estimate how many steps of unanchored strategic reasoning individuals carry out. The notation for *s* steps of unanchored strategic reasoning is *s*_{unanch}.

The following text formally describes unanchored strategic reasoning. Let \mathcal{D}_{jg} be the family of all strategies (i.e., distributions) supported on $[\alpha_{Lg}(j), \alpha_{Ug}(j)]$. Then define

$$\tilde{\Sigma}_{jg}^0 = \{\sigma_j \in \mathcal{D}_{jg}\}$$

Similarly, define

$$\Sigma_{jg}^0 = \left[\alpha_{Lg}(j), \alpha_{Ug}(j) \right]$$

to be the set of actions that are consistent with the use of zero steps of unanchored strategic reasoning by player *j* in game *g*. Of course, by construction, Σ_{jg}^{0} is the entire

⁷The model could be extended to games with multiple Nash equilibria as long as all games under study have the same number of Nash equilibria, and those equilibria can be distinguished according to some criterion. For example, suppose the games under study have two Nash equilibria that can be distinguished in some observable way. Potentially, for example, one Nash equilibrium could be focal while the other equilibrium is nonfocal, or one Nash equilibrium could satisfy a certain equilibrium refinement, while the other equilibrium does not satisfy that equilibrium refinement, or, more simply, one Nash equilibrium could result in a larger choice of action, whereas the other equilibrium results in a smaller choice of action. Other distinguishing properties are also possible. Then those two Nash equilibria could both be included in the model in the same way that different numbers of steps of anchored strategic reasoning are included in the model, and the identification strategy could be used to identify that model. The same considerations would apply if a certain number of steps of anchored strategic reasoning predicted two actions.

⁸See, for example, Tan and da Costa Werlang (1988) or Fudenberg and Tirole (1991). Level-k rationality has been *assumed* to generate the data in Aradillas-Lopez and Tamer (2008), Kline and Tamer (2012), and Kline (2015) to identify the utility function.

action space. Then, for $s \ge 0$, define

$$\tilde{\Sigma}_{jg}^{s+1} = \left\{ \sigma_j \in \mathcal{D}_{jg} : \exists \sigma_{-j} \in \prod_{j' \neq j} \operatorname{co}(\tilde{\Sigma}_{j'g}^s) \text{ s.t. } \sigma_j \text{ is supported on} \right.$$

the set of solutions to
$$\max_{a_j \in [\alpha_{Lg}(j), \alpha_{Ug}(j)]} E_{\sigma_{-j}}(u_{jg}(a_1, \dots, a_{M_g})) \right\}.$$

These are the strategies σ_j for which there are strategies σ_{-j} of the other players that can be used by other players who use strategies in $\prod_{j'\neq j} \operatorname{co}(\tilde{\Sigma}_{j'g}^s)^9$ such that σ_j is the best response to the other players using those strategies. Similarly, define

$$\Sigma_{jg}^{s+1} = \left\{ a_j \in \left[\alpha_{Lg}(j), \alpha_{Ug}(j) \right] : \exists \sigma_{-j} \in \prod_{j' \neq j} \operatorname{co}(\tilde{\Sigma}_{j'g}^s) \text{ s.t.} \right.$$
$$a_j \in \operatorname*{arg\,max}_{a_j \in [\alpha_{Lg}(j), \alpha_{Ug}(j)]} E_{\sigma_{-j}} \left(u_{jg}(a_1, \dots, a_{M_g}) \right) \right\}$$

to be the set of actions that are consistent with the use of s + 1 steps of unanchored strategic reasoning by player *j* in game *g*.

Note the intuitive appeal of *s* steps of unanchored strategic reasoning in terms of iterated deletion of dominated strategies, especially in the case of two-player games. Intuitively, strategies in $\tilde{\Sigma}_{jg}^1$ are best responses to some strategies of the opponents and, therefore, survive one round of deletion of dominated strategies; $\tilde{\Sigma}_{jg}^2$ are best responses to some strategies of the opponents that survive one round of deletion of dominated strategies; $\tilde{\Sigma}_{jg}^2$ are best responses to some strategies and, therefore, survive two rounds of deletion of dominated strategies, and so forth. See Tan and da Costa Werlang (1988) or Fudenberg and Tirole (1991) for further details.

After finding the set Σ_{jg}^s , player *j* in game *g* that uses *s* steps of unanchored strategic reasoning must use some selection rule to select an action to actually play from the set Σ_{jg}^s . By definition, a selection rule is a distribution supported on Σ_{jg}^s . The *consistency with the application of iterated deletion of dominated strategies* does not place any further restrictions on the selection rule. However, the definition of using *unanchored strategic reasoning* entails the use of a specific selection rule. Let $\psi_{jg}(\cdot)$ be a known strictly positive and continuous function on the action space $[\alpha_{Lg}(j), \alpha_{Ug}(j)]$. Then suppose that player *j* in game *g* that uses *s* steps of unanchored strategic reasoning takes an action $a \in \Sigma_{jg}^s$ with "density"

$$\zeta_{jg}^{s}(a) = \frac{\psi_{jg}(a)}{\Psi_{jg}(\Sigma_{jg}^{s})},$$

where $\Psi_{jg}(\Sigma_{jg}^s) = \int_{\Sigma_{jg}^s} \psi_{jg}(a) d\mu(a; \Sigma_{jg}^s)$ and $\mu(\cdot; \Sigma_{jg}^s)$ is the appropriate dominating measure for a distribution on Σ_{jg}^s .¹⁰

⁹Due to use of the convex hull operator co, this allows mixtures of strategies, in cases of nonconvexity.

¹⁰Therefore, this implicitly requires that such a distribution exists. Consequently, it is implicitly assumed that Σ_{ig}^{s} is either Lebesgue measurable with nonzero and finite measure or is a finite set. In particular, if

An important special case is $\psi_{jg}(\cdot) \equiv 1$, so that the selection rule $\zeta_{jg}^s(\cdot)$ is the uniform distribution on Σ_{jg}^s . By construction, all actions in Σ_{jg}^s are equally consistent with using *s* steps of unanchored strategic reasoning. Equivalently, there is equal justification for using each of the actions in Σ_{jg}^s based on using *s* steps of unanchored strategic reasoning. Consequently, assuming that actions are played with probability proportionate to the justification for playing that action, similar to the principle of indifference, the selection rule associated with *s* steps of unanchored strategic reasoning would indeed be the uniform distribution over Σ_{jg}^{s} .¹¹ The uniform distribution appears consistent with the data in the empirical application, as discussed in Section 5.2. More generally, the selection rule can be biased toward or against the use of an action *a* proportionate to the quantity $\psi_{jg}(a)$. For example, this accommodates situations where the econometrician knows that certain actions are focal. Specifically, $\psi_{jg}(\cdot)$ can be a unimodal density, so $\psi_{jg}(\cdot)$ is large around the mode (i.e., the focal action) and small away from the mode (i.e., the nonfocal actions).

This selection rule is such that a player who uses *s* steps of unanchored strategic reasoning does not always use an action consistent with refinements of *s* steps of unanchored strategic reasoning, since the density is strictly positive on all of Σ_{jg}^s . Therefore, the selection rule guarantees that the use of *s* steps of unanchored strategic reasoning has a distinct definition from the use of refinements of *s* steps of unanchored strategic reasoning. Otherwise, if the use of two decision rules cannot be distinguished even by definition, then the use of those two decision rules could never be distinguished using data, resulting in a failure by definition of point identification.

For example, the Nash equilibrium action is also consistent with *s* steps of unanchored strategic reasoning, that is, $c_{jg}(NE) \in \Sigma_{jg}^s$. If there were no restrictions on the selection rule, then the selection rule could be that a player who uses *s* steps of unanchored strategic reasoning always uses the Nash equilibrium action. If so, then it would be impossible by definition to distinguish between the use of the Nash equilibrium and the use of *s* steps of unanchored strategic reasoning. Similarly, any action consistent with *s'* steps of unanchored strategic reasoning is also consistent with $0 \le s \le s'$ steps of unanchored strategic reasoning is also consistent with $0 \le s \le s'$ steps of unanchored strategic reasoning is also consistent with $0 \le s \le s'$ steps of unanchored strategic reasoning is also consistent with $0 \le s \le s'$ steps of unanchored strategic reasoning is also consistent with $0 \le s \le s'$ steps of unanchored strategic reasoning is also consistent with $0 \le s \le s'$ steps of unanchored strategic reasoning is also consistent with $0 \le s \le s'$ steps of unanchored strategic reasoning is also consistent with $0 \le s \le s'$ steps of unanchored strategic reasoning, that is, $\Sigma_{jg}^{s'} \subseteq \Sigma_{jg}^{s}$ for $0 \le s \le s'$. Therefore, an action can be consistent with many different numbers of steps of unanchored strategic reasoning, making it difficult to infer the number of steps of unanchored strategic reasoning used to generate that action. If there were no restrictions on the selection rule, then the selection rule could be that a player who uses *s* steps of unanchored strategic reasoning always uses an action consistent with s + 1 steps (or some other greater number of steps)

 $[\]Sigma_{jg}^s$ is Lebesgue measurable with nonzero and finite measure, as in the empirical application where Σ_{jg}^s are intervals, then $\mu(\cdot; \Sigma_{jg}^s)$ is Lebesgue measure and $\zeta_{jg}^s(\cdot)$ is an ordinary density; if Σ_{jg}^s is a finite set, then $\mu(\cdot; \Sigma_{jg}^s)$ is counting measure and $\zeta_{jg}^s(\cdot)$ is a "density" with respect to counting measure, otherwise known as a probability mass function. It is implicitly assumed that indeed these integrals exist and are finite under the condition that $\psi_{jg}(\cdot)$ is integrable with respect to the appropriate dominating measure(s).

¹¹One statement of the principle of indifference from Carnap (1953, p. 193) is "[i]f no reasons are known which would favor one of several possible events, then the events are to be taken as equally probable." Similarly, the interpretation of the uniform selection rule is that if no reasons are known that would favor any of the actions within Σ_{ie}^{s} , then those actions are to be taken with equal probability.

of unanchored strategic reasoning. If so, then it would be impossible by definition to distinguish between the use of s and s + 1 steps of unanchored strategic reasoning.

Moreover, this selection rule is such that the "relative bias" between taking actions $a^{(1)}$ and $a^{(2)}$ consistent with *s* steps of unanchored strategic reasoning is the same for all *s*, since that relative bias $\frac{\zeta_{jg}^{s}(a^{(1)})}{\zeta_{jg}^{s}(a^{(2)})} = \frac{\psi_{jg}(a^{(1)})}{\psi_{jg}(a^{(2)})}$ does not depend on *s* as long as $a^{(1)}$ and $a^{(2)}$ are indeed consistent with *s* steps of unanchored strategic reasoning. This makes it possible to distinguish between the use of a certain number *s* steps of unanchored strategic reasoning and the use of a different number *s'* steps of unanchored strategic reasoning with a selection rule that is biased toward taking actions consistent with *s* steps of unanchored strategic reasoning.

Further, it is possible to treat the selection rule on unanchored strategic reasoning as another parameter in the model. Under suitable restrictions on the class of admissible selection rules, which are equivalent to restrictions on the class of admissible $\psi_{jg}(\cdot)$ functions, Appendix C discusses identification of the selection rule. Hence, there are two possible interrelated approaches to the selection rule: (a) the econometrician can define using unanchored strategic reasoning to involve the uniform selection rule or some other known selection rule and directly apply the identification result, or (b) the econometrician can expand the model to allow that using unanchored strategic reasoning involves the use of some unknown selection rule within a suitably restricted class of selection rules, identify the selection rule per Appendix C, and then apply the identification result with that identified (known from the data) selection rule. By either varying the "known" selection rule or expanding the model and treating the selection rule as another parameter, this implies the ability to conduct sensitivity analysis with respect to the selection rule part of the definition of using unanchored strategic reasoning.

REMARK 2.1 (Epistemic Interpretation). The results of Tan and da Costa Werlang (1988) can be used to provide an epistemic interpretation of the set of strategies $\tilde{\Sigma}_{jg}^s$. For s = 1, using a strategy in $\tilde{\Sigma}_{jg}^s$ is equivalent to being rational (at least), and for $s \ge 2$, using a strategy in $\tilde{\Sigma}_{jg}^s$ is equivalent to being rational and also knowing everyone (knows

¹²Using the formal notation introduced subsequently in the paper, this is an implication of the more general fact that arbitrary mixtures of densities are not point identified. For example, suppose that there is only one strategic behavior type (i.e., R = 1), and suppose that the econometrician assumes that strategic behavior type uses either 0 or 1 steps of unanchored strategic reasoning (i.e., $\mathcal{U} = \{0_{\text{unanch}}, 1_{\text{unanch}}\}$ and $\mathcal{A} = \mathcal{M} = \emptyset$). Then a specification of the model would entail, for that one type, the specification of $\Lambda_1(0_{\text{unanch}})$ and $\Lambda_1(1_{\text{unanch}})$, and also, for each game, a distribution H_{g0} that is supported on the actions associated with 0 steps of unanchored strategic reasoning in game g, and a distribution H_{1g} that is supported on the actions associated with 1 step of unanchored strategic reasoning in game g. Without any restrictions on the selection rule, H_{0g} and H_{1g} could be any distributions with the appropriate support. The specification $(\Lambda_1(0_{\text{unanch}}), \Lambda_1(1_{\text{unanch}}), H_{0g}, H_{1g})$ implies the observed distribution of actions in game g that is given by the mixture $\Lambda_1(0_{\text{unanch}})H_{0g} + \Lambda_1(1_{\text{unanch}})H_{1g}$. Now consider the strategic behavior type that uses 0 steps of unanchored strategic reasoning with probability 1, and uses the $\Lambda_1(0_{\text{unanch}})H_{0g} + \Lambda_1(1_{\text{unanch}})H_{1g}$ distribution on Σ_{1g}^0 . By construction, that results in the same observed distribution of actions. Intuitively, this can happen if individuals who use 0 steps of unanchored strategic reasoning are biased toward using the actions that are also consistent with using 1 step of unanchored strategic reasoning. Consequently, these two specifications of the model are observationally equivalent.

everyone)^{*s*-2} is rational (at least), in addition to some other conditions including those related to players acting independently of each other. Further, *rationalizability*, defined as using a strategy in the set $\bigcap_{s=1}^{\infty} \tilde{\Sigma}_{jg}^{s}$, is roughly equivalent to common knowledge of rationality in addition to some other conditions, including those related to players acting independently of each other.

For example, in a two-player game, player 1 who uses a strategy in $\tilde{\Sigma}_{1g}^2$ can be interpreted "as if" to use the following strategic reasoning: I think my opponent will use strategy σ_2 . I think my opponent will use σ_2 because σ_2 would be a best response from the perspective of my opponent if I were to use strategy σ_1 . And given that I think my opponent will use σ_2 , I should use the strategy σ_1' as a best response to σ_2 .

REMARK 2.2 (Consistency With Iterated Deletion of Dominated Strategies). The experimental game theory literature has sometimes checked whether observed actions are consistent with certain solution concepts or decision rules—in particular the steps of iterated deletion of dominated strategies—as a stand-alone exercise separate from, for example, estimating a structural level-k model. See, for example, the discussion in footnote 20 in the context of Costa-Gomes and Crawford (2006). In contrast, in this paper, unanchored strategic reasoning is included as a decision rule in a model alongside other decision rules, making it possible to answer the question of how often (and/or whether) subjects use a given number of steps of unanchored strategic reasoning. Note the fundamental distinction between consistency with and actually using a given number of steps of unanchored strategic reasoning. An action can be consistent with a given number of steps of unanchored strategic reasoning even though the player taking that action did not use that number of steps of unanchored strategic reasoning. For example, a player might use two steps of unanchored strategic reasoning, but nevertheless take an action that is also consistent with both zero and one steps of unanchored strategic reasoning. The fact that an action is consistent with a given number of steps of unanchored strategic reasoning is not necessarily evidence that the player taking that action actually used that number of steps of unanchored strategic reasoning.

2.2.3 Anchored strategic reasoning It is possible to add to the above iterated definitions the condition that, for all players j and games g, $\tilde{\Sigma}_{jg}^0$ consists of only one strategy: the uniform distribution over the action space. This results in anchored strategic reasoning, because the steps of strategic reasoning become anchored to the uniform distribution being used by players who use zero steps of strategic reasoning. In the experimental game theory literature, with citations provided in the Introduction, this is known as the level-k model, but the terms "anchored" and "unanchored" are used in this paper to emphasize the relationship between the two classes of decision rules. The notation for s steps of anchored strategic reasoning is s_{anch} .

Zero steps of unanchored strategic reasoning are observationally equivalent to zero steps of anchored strategic reasoning, at least under the condition of a uniform selection rule on unanchored strategic reasoning, but anchoring does revise the implications of using more than zero steps of strategic reasoning by working through the iterated definition of steps of strategic reasoning described in Section 2.2.2. For example, a player

who uses one step of anchored strategic reasoning would use a strategy that is a best response to the other players using the strategy that is the uniform distribution over the action space, and a player who uses two steps of anchored strategic reasoning would use a strategy that is the best response to the other players using a strategy consistent with one step of anchored strategic reasoning.

The results are derived based on the assumption that there is a unique action consistent with anchored strategic reasoning (for each $s \ge 1$), as is typically the case for games studied in the related experimental game theory literature: player j in game g who uses s steps of anchored strategic reasoning takes action $c_{jg}(s_{anch})$. There is typically a range of actions consistent with s steps of unanchored strategic reasoning. Hence, it is possible to distinguish an individual who uses unanchored strategic reasoning from an individual who uses anchored strategic reasoning, because the latter will *always* take the action associated with anchored strategic reasoning, whereas the former will not.

2.2.4 *Assumptions on strategic reasoning* Assumption 2.1 states that the set of steps of strategic reasoning that subjects might use is known by the econometrician to be a finite set. This is consistent with prior experimental results, which indicate individuals use a very small number of steps of reasoning. The consequence of Assumption 2.1 is that subjects are restricted to using a finite set of decision rules, rather than an infinite set of decision rules. It would be extremely difficult to distinguish between infinitely many decision rules, especially with finite data.

ASSUMPTION 2.1 (Steps of Strategic Reasoning). The numbers of steps of unanchored strategic reasoning that subjects might use is the known finite set U. The numbers of steps of anchored strategic reasoning that subjects might use is the known finite set A.

2.3 Computational mistakes

Roughly following the literature on experimental game theory, computational mistakes arise when a subject "intends" to use a certain decision rule, but fails to correctly take the associated action. The decision rules subject to computational mistakes are the decision rules that are associated with a *unique* action, collected in the set \mathcal{M} : the steps of anchored strategic reasoning and Nash equilibrium.¹³ The econometrician can assume ex ante that subjects do not make computational mistakes, in which case the sufficient conditions for point identification are weaker.

Let $\xi(\cdot)$ be a known bounded and continuous density defined on support [-1, 1] that is bounded away from zero in the sense that $\xi(x) \ge \kappa > 0$ for all $x \in [-1, 1]$ for some κ . Conversely, $\xi(\cdot)$ is zero off the support [-1, 1], by definition of support. The continuity at the endpoints -1 and 1 is implicitly understood to be right and left continuity. Suppose that subject *i* "intends" to use a particular decision rule in \mathcal{M} that predicts the action *c*, and that subject *i* is playing the role of player *j* in game *g*. There is δ_i probability that the

¹³Computational mistakes arise only with decision rules that are associated with a unique action (which is where computational mistakes have been allowed in the prior literature), avoiding the ambiguity about what it would mean to incorrectly compute the action associated with a decision rule that is consistent with a range of actions, as in unanchored strategic reasoning.

subject makes a computational mistake. If there is a computational mistake, then the subject actually takes an action according to the $\xi(\cdot)$ density, translated to an interval of radius $\rho_i(\alpha_{Ug}(j) - \alpha_{Lg}(j))$ that is centered at the "intended" action *c*, intersected with the action space $[\alpha_{Lg}(j), \alpha_{Ug}(j)]$,

$$[\alpha_{Lg}(j), \alpha_{Ug}(j)] \cap [c - \rho_i (\alpha_{Ug}(j) - \alpha_{Lg}(j)), c + \rho_i (\alpha_{Ug}(j) - \alpha_{Lg}(j))]$$

=
$$[\max\{\alpha_{Lg}(j), c - \rho_i \Omega_{jg}\}, \min\{\alpha_{Ug}(j), c + \rho_i \Omega_{jg}\}],$$

with $\Omega_{jg} \equiv \alpha_{Ug}(j) - \alpha_{Lg}(j)$. The intersection with $[\alpha_{Lg}(j), \alpha_{Ug}(j)]$ guarantees that the action is within the action space. Consequently, the subject takes an action *a* according to the density

$$\omega_{jg,c,\rho_i}(a) = \frac{2 \times \xi \left(\frac{a - \frac{\min\left\{\alpha_{Ug}(j), c + \rho_i \Omega_{jg}\right\} + \max\left\{\alpha_{Lg}(j), c - \rho_i \Omega_{jg}\right\}}{2}}{\frac{\min\left\{\alpha_{Ug}(j), c + \rho_i \Omega_{jg}\right\} - \max\left\{\alpha_{Lg}(j), c - \rho_i \Omega_{jg}\right\}}{2}}\right)}{\min\left\{\alpha_{Ug}(j), c + \rho_i \Omega_{jg}\right\} - \max\left\{\alpha_{Lg}(j), c - \rho_i \Omega_{jg}\right\}}.$$

The parameter ρ_i characterizes the magnitude of computational mistakes: larger ρ_i implies the possibility of larger computational mistakes. The range of computational mistakes is ρ_i multiplied by the width of the action space $\Omega_{jg} \equiv \alpha_{Ug}(j) - \alpha_{Lg}(j)$ to reflect the fact that games with larger action spaces are more subject to relatively larger computational mistakes. The model of computational mistakes is formalized in Assumption 2.2. Similar identification strategies could be used for similar models of computational mistakes.

ASSUMPTION 2.2 (Computational Mistakes). Either of the following scenarios holds:

(i) The econometrician allows the possibility of computational mistakes. The probability that subject *i* makes a computational mistake is $0 \le \delta_i < 1$. The magnitude of the computational mistakes made by subject *i* is $\rho_i > 0$. If subject *i* makes a computational mistake in game *g* as player *j* and intended to use a decision rule that would result in taking action *c*, then subject *i* takes an action according to the $\xi(\cdot)$ density, translated to $[\alpha_{Lg}(j), \alpha_{Ug}(j)] \cap [c - \rho_i(\alpha_{Ug}(j) - \alpha_{Lg}(j)), c + \rho_i(\alpha_{Ug}(j) - \alpha_{Lg}(j))]$. The econometrician knows $\overline{\rho}$ such that $\rho_i < \overline{\rho}$ for all subjects *i*.

(ii) The econometrician does not allow the possibility of computational mistakes and, therefore, knows that $\delta_i \equiv 0$ and $\rho_i \equiv 0$ for all subjects *i*. For the purposes of future assumptions, the econometrician sets $\overline{\rho} = 0$.

If the econometrician allows the possibility of computational mistakes, then it is assumed that $\rho_i > 0$ for all subjects *i*. If $\rho_i = 0$ were allowed, then there would be a complication relating to the fact that subjects who do not make computational mistakes ($\delta_i = 0$) are observationally equivalent to subjects who do make computational mistakes with zero magnitude ($\delta_i > 0$ but $\rho_i = 0$).

2.4 Strategic behavior rules and within-individual heterogeneity

Each subject i has a strategic behavior rule

$$\theta_i = \left(\lambda_i(\cdot), \delta_i, \rho_i\right)$$

that characterizes how it behaves in games. These are ex ante unknown by the econometrician. The components of the strategic behavior rule are as follows:

(i) The distribution $\lambda_i(\cdot)$ over decision rules characterizes the probabilities that subject *i* uses each decision rule. The argument of $\lambda_i(\cdot)$ is a decision rule. For example, $\lambda_i(\text{NE})$ is the probability that subject *i* uses the Nash equilibrium (NE) solution concept when it plays a game. The decision rules are described in Section 2.2.

(ii) The parameters δ_i and ρ_i are the probability and magnitude of computational mistakes made by subject *i*. As described in Section 2.3, a subject might "intend" to use a particular decision rule, but fail to compute the associated action correctly and actually take an action that is only approximately equal to the action predicted by the intended decision rule. A special case of the model rules out computational mistakes.

The distribution over decision rules allows the existence of within-individual heterogeneity: a given subject might use multiple decision rules. However, the model does not impose the existence of within-individual heterogeneity, thereby nesting related models in which subjects each use only one decision rule as special cases. Indeed, the empirical application finds evidence of within-individual heterogeneity without ex ante imposing the existence of within-individual heterogeneity. If the econometrician restricts the parameter space for $\lambda_i(\cdot)$ to be degenerate distributions that place probability 1 on just one decision rule, then the econometrician assumes away within-individual heterogeneity.

The distribution over decision rules is an exogenous and fixed characteristic of a subject. Therefore, the model does not allow the possibility that subjects endogenously adjust their probabilities of using the decision rules due to learning or related dynamic considerations. Although there is an important literature on learning in games (e.g., as discussed in Camerer (2003, Chapter 6)), the model in this paper follows a significant part of the experimental game theory literature that abstracts from learning. Experiments, including the experiment analyzed in the empirical application in Section 5, may be designed specifically to reduce or even eliminate the possibility of learning and related dynamic considerations. The experiment can limit the feedback presented to the subjects about their play of the games until the completion of the entire experiment. Further, the experiment can begin with an initial learning period that is not analyzed by the econometrician, so that the data that are analyzed by the econometrician are subsequent to the subjects learning how to play the game. Similarly, the model does not allow the possibility that the games played in the experiment have different difficulties that would lead to the use of different decision rules.

Therefore, if an experiment involves learning and/or different difficulties of the games, then the model is misspecified relative to that experimental data. Because these features of the experiment would result in individuals using multiple decision rules over the course of the experimental study, estimates of the model based on such data could

be expected to "account for" those features of the experimental data by estimating that individuals have within-individual heterogeneity, since within-individual heterogeneity does result in individuals using multiple decision rules. Hence, within-individual heterogeneity could "fit" the data coming from an experiment involving learning and/or different difficulties of the games, in the usual sense that estimating a misspecified model results in a "best fit" of the data relative to the misspecified model. Therefore, the interpretation of estimates from the model that indicate the existence of withinindividual heterogeneity depend on the credibility of the assumption that learning and/or different difficulties of the games are not features of the experiment. These features of the experiment would also be problematic for models that abstract from withinindividual heterogeneity, because those models assume that each individual always uses the same decision rule. Another reason that individuals might not always conform to the predictions of one decision rule is that they make computational mistakes, as accommodated in the model and discussed in Section 2.3.

In single-agent decision problem experiments, there is evidence that individuals do not always make the same choices when repeatedly faced with the same decision problem (e.g., Rieskamp, Busemeyer, and Mellers (2006) and Rieskamp (2008)). One explanation for such observed behavior is random utility,¹⁴ because as Machina (1985, p. 575) notes, "if when confronted with a choice over two objects the individual chooses each alternative a positive proportion of the time, it seems natural to suppose that this is because he or she 'prefers' each one to the other those same proportions of the time." Bardsley, Cubitt, Loomes, Moffatt, Starmer, and Sugden (2010, Section 7.2.3) summarize random utility models to have the structure "(i) that the individual's preferences can be represented by some set of functions, all of which are consistent with that theory; and (ii) that for any particular decision task, the individual acts as if she picks one of those functions at random from the set and applies it to the task in question; then (iii) 'puts back' that function into the set before picking again at random when tackling another decision (even if it is the identical task encountered another time)." McFadden (1981, p. 205) summarizes one common interpretation of random utility models as "[t]hen the individual is a classical utility maximizer given his state of mind, but his state of mind varies randomly from one choice situation to the next."

Similar to how random utility models allow that the behavior of individuals in singleagent decision problems may be described as arising from randomly selecting from a set of utility functions, within-individual heterogeneity allows that the behavior of individuals in games may be described as arising from randomly selecting from a set of decision rules. Further, it could be that the behavior of individuals in games may be described as arising from randomly selecting from a set of beliefs about the type of their opponent and using the induced decision rule that best responds to that belief. So, the "state of mind" is the belief about the type of their opponent. For example, in the level-*k* model of thinking (i.e., anchored strategic reasoning in this paper, detailed in Section 2.2), an

¹⁴Many models in econometrics known as random utility tend to be interpreted to emphasize a distribution of utility across the population. In contrast, these "random utility" models emphasize a distribution of utility for a given agent. The difference essentially concerns whether the "randomness" is a fixed characteristic of individual agents across decision problems.

individual who believes the opponent is level 0 with probability p and level-1 with probability 1 - p will use the level-1 strategy with probability p and the level-2 strategy with probability 1 - p. This differs from the standard approach to responding to uncertainty about the type of the opponent, which would not generate within-individual heterogeneity, because it would entail individuals using the strategy that is the best response to the entire distribution of beliefs about the type of the opponent. By resolving uncertainty about the opponent *before* taking an action, individuals can exhibit within-individual heterogeneity.

Therefore, different decision rules may be interpreted as involving different beliefs about the strategy of the opponent, and beliefs about the strategy of the opponent affect the (expected) utility each individual associates with each of its actions. Hence, randomly selecting from a set of decision rules is related to randomly selecting from a set of utility functions, where the utility functions in the set of utility functions differ because of the different beliefs held about the strategy of the opponent. Consequently, randomly selecting from a set of decision rules provides a justification for why individuals might appear as if to randomly select from a set of utility functions, as in the standard formulation of random utility models.

More generally, especially in the empirical experimental game theory literature concerning the level-*k* model of thinking (i.e., anchored strategic reasoning in this paper), issues related to but distinct from within-individual heterogeneity have been investigated as a sort of robustness check on the stability of the estimates. The details vary across papers, but two main questions are common. See Stahl and Wilson (1995) or Georganas, Healy, and Weber (2015) for some examples. One question concerns checking whether the aggregate distribution of behavior (i.e., the fraction of level-1 behavior, the fraction of level-2 behavior, etc.), which is the same as the aggregate distribution of types (i.e., the fraction of level-1 thinkers, the fraction of level-2 thinkers, etc.) in models that assume that each individual exclusively uses one decision rule, appears to be the same across multiple sets of games. The example in Section 3.2 shows that the fraction of individuals exhibiting any given number of steps of reasoning can be the same across games, even though particular individuals do exhibit within-individual heterogeneity. Therefore, questions concerning the aggregate distribution of behavior are distinct from questions concerning within-individual heterogeneity, and indeed withinindividual heterogeneity can be obscured when investigating only the aggregate distribution of behavior, because within-individual heterogeneity is a characteristic of an individual, not aggregate behavior across individuals. Another question concerns checking whether a particular individual is estimated to be the same type across multiple sets of games (or, more or less equivalently, whether individuals appear to statistically conform out of sample to their estimated type). This question is more similar to, but still distinct from, questions concerning within-individual heterogeneity. An individual who most often uses a particular decision rule is likely to always be estimated to be the type who uses that decision rule, across different sets of games, since that provides the best fit among the types restricted to using one decision rule, regardless of underlying within-individual heterogeneity. More generally, models that are restricted to estimating each individual to be a type that exclusively uses one decision rule (e.g.,

level 1 or level 2) are misspecified in the presence of within-individual heterogeneity. In contrast, the model in this paper explicitly allows within-individual heterogeneity.

Although the above discussion describes possible explanations for within-individual heterogeneity, further research would be needed to understand and distinguish between the possible sources of within-individual heterogeneity. Within-individual heterogeneity is a characteristic of observed behavior, with potentially many "as if" explanations. Consequently, the model provides a framework for studying the observable implications of within-individual heterogeneity, but does not provide an explanation for why individuals do or do not exhibit within-individual heterogeneity. Similarly, related papers provide frameworks for studying the observable implications of an economic theory (e.g., Nash equilibrium, level-*k* thinking, etc.), without attempting to explain why individuals do or do not conform to the predictions of that theory.

2.5 Strategic behavior types

The model is based on the condition that there are at most *R* strategic behavior rules used in the population. Hence, by definition, there are at most *R* strategic behavior types, indexed by r = 1, 2, ..., R, and denoted by $\Theta_r = (\Lambda_r, \Delta_r, P_r)$, where the quantities comprising Θ_r in uppercase letters correspond to the quantities comprising θ_i in lowercase letters. Therefore, for strategic behavior type *r*, Λ_r is the distribution over decision rules, Δ_r is the probability of a computational mistake, and P_r is the magnitude of a computational mistake.

The population fraction of subjects who are type r is $\pi(r)$. It is allowed that $\pi(r) = 0$ for some r, so that fewer than R strategic behavior types exist. When subject i is born, it is assigned to use strategic behavior type $\Theta_{\tau(i)}$, where $\tau(i) \in \{1, 2, ..., R\}$, according to the distribution $\pi(\cdot)$ over $\{1, 2, ..., R\}$. Therefore, by construction, $\theta_i = \Theta_{\tau(i)}$. The condition that the population uses at most R strategic behavior rules guarantees that the model is parsimonious and is required for the identification strategy.

Although *R* is known by the econometrician, $\{\Theta_r, \pi(r)\}_{r=1}^R$ are unknown by the econometrician. Consequently, the econometrician knows that there are at most *R* strategic behavior rules used in the population, but the econometrician does not know those strategic behavior rules or the population fractions of subjects that use each strategic behavior rule. Indeed, the identification result shows sufficient conditions for point identification (and therefore estimation) of the unknown $\{\Theta_r, \pi(r)\}_{r=1}^R$.

2.6 Behavioral implications of the model

The behavioral implications of the model can be described in the following procedural way.

I. Each subject is born and permanently assigned its strategic behavior rule $\theta_i = (\lambda_i(\cdot), \delta_i, \rho_i)$ by nature per Section 2.5.

II. Each time subject *i* encounters a game to play, the following events occur:

158 Brendan Kline

(a) Subject *i* chooses the decision rule it intends to use in that game. The probability that subject *i* chooses decision rule *k* is $\lambda_i(k)$. The subject might choose, for example, to use the Nash equilibrium or to use one step of unanchored strategic reasoning. The set of decision rules is described in Section 2.2.

(b) If the intended decision rule is not subject to computational mistakes, as described in Section 2.3, then the subject takes an action according to that decision rule. Otherwise, the subject attempts to compute the action associated with the intended decision rule. The subject either correctly or incorrectly computes the action:

(i) The probability of correct computation is $1 - \delta_i$. In this case, the subject actually takes the action associated with the intended decision rule.

(ii) The probability of incorrect computation is δ_i . In this case, the subject actually takes an action that is only approximately equal to the action associated with the intended decision rule. The details of computational mistakes are described in Section 2.3.

For example, if $\lambda_i(k) = 0.2$ and $\delta_i = 0.05$, then subject *i* uses decision rule *k* with probability 0.2. Supposing that *k* is subject to computational mistakes, with probability 0.95, it correctly computes the decision rule and actually does take the associated action, but, with probability 0.05, it makes a small computational mistake and takes an action that is only approximately equal to the action associated with decision rule *k*.

2.7 Data and sketch of identification problem

The data observed by the econometrician are the actions taken by each of N subjects in games of the sort described in Section 2.1. The subjects are indexed by i = 1, 2, ..., N. Each subject plays each of G games, indexed by g = 1, 2, ..., G. It is assumed essentially without loss of generality, by redefining the player roles appropriately, that the subjects in the data set are always player 1 in the games.¹⁵ The observed action of subject i in game g is y_{ig} . See the empirical application in Section 5 for one of many instances of such a data set from the experimental game theory literature. As discussed in Section 2.4, because of the non-equilibrium nature of the analysis, the actions of the opponents of a subject are not relevant, since the analysis focuses on identifying/estimating the solution concept(s) or decision rule(s) that generate the behavior of individual subjects. In

¹⁵In many experiments, including the experiment in the empirical application, for each actual game in the experiment, each subject plays in all player roles in that game. Then it is possible to redefine the games to satisfy the condition that the subjects in the data set are always player 1 in the games. For example, if there is a game with a row player and a column player, then the data from the subjects' behavior in the row player role of the game can be "game 1" data and the data from the subjects' behavior in the column player role of the game can be "game 2" data. Even if the experimental design does not intentionally assign all subjects to play in all player roles, as long as the assignment to player roles is suitably randomized and exogenous to the model, there can be some representative subset of subjects who do happen to always play as player 1 in all games, and the identification results can be viewed as showing sufficient conditions for achieving identification of the model parameters based on that subset of the subjects who play as player 1 in all games. Of course, the other data not from that subset of subjects would also have identifying content and would be used in estimation.

principle, it would be enough for an experiment to present each subject with each of the games, without presenting the games to the opponents.

The population distribution of the observed data is $P(\{y_g\}_{g=1}^G)$, the distribution of actions in the *G* games across the population of subjects. The identification problem is to establish sufficient conditions under which it is possible to uniquely recover the unknown parameters $\{\Theta_r, \pi(r)\}_{r=1}^R$ from $P(\{y_g\}_{g=1}^G)$.

The identification problem corresponds to $N \to \infty$ while *G* is fixed, corresponding to data from a population of subjects observed to play *G* games. An alternative identification problem corresponds to $N \to \infty$ and $G \to \infty$, corresponding to data from a population of subjects observed to play a population of games. Identification in that setup would require making assumptions about the population distribution of games, including games that are not among the finitely many games in the experiment. It seems difficult to interpret assumptions on games that are not in the experiment. In contrast, the identification results in the fixed *G* setup require only that the econometrician verify that the games in the experiment satisfy certain conditions. Therefore, identification results in the $G \to \infty$ setup could be less persuasive than identification results in the fixed *G* setup.

The discussion focuses on the experimental design involving N subjects, each of whom plays each of G games once. Another experimental design involves just one subject who repeatedly plays one game. Assuming away the learning and other dynamic considerations discussed in Section 2.4, the model of one subject playing one game repeatedly, viewing each play of the game as an observation indexed by i = 1, 2, ..., N, is mathematically equivalent to many subjects of the same strategic behavior type (i.e., R = 1) playing one game once each, viewing each subject's play of the game as an observation indexed by i = 1, 2, ..., N. In both cases, there is a single strategic behavior type that generates the data, whether the data happen to be generated by one subject repeatedly playing the game (necessarily described by one strategic behavior type) or many subjects playing the game (all described by the same strategic behavior type). Therefore, the identification problem arising from this alternative experimental design is covered by the identification results in this paper, taking N to be the number of times the subject plays the game, G = 1 to reflect that the subject plays one game, and R = 1 to reflect that the subject is necessarily of one strategic behavior type. In particular, as discussed in Section 3.3, the identification problem remains challenging even in this special case of the model.

As with estimation of any model, the empirical results are necessarily relative to the environment that generates the data. Data sets from different environments might result in different estimates of the model parameters, because of differences in the environments. For example, estimation of the model on different populations of individuals might reveal that different populations are composed of different types, or composed of the same types in different proportions. Similarly, estimation of the model on different sets of games (e.g., different difficulties of the games) might reveal that different sets of games result in the use of different decision rules or the same decision rules with different probabilities.

3. Setup of the identification problem

The identification problem in this model concerns the question of whether it is possible to recover the parameters of the model (i.e., $\{\Theta_r, \pi(r)\}_{r=1}^R$) from the population distribution of the data. The model will fail to be point identified if it happens that more than one specification of the parameters generates the same distribution of the data, because then the "true" specification of the parameters cannot be distinguished from a "false" specification of the parameters. Therefore, point identification is a prerequisite for estimating the parameters of the model.¹⁶ The parameters of the model are not point identified without nontrivial sufficient conditions, as Section 3.2 provides a counterexample to point identification in the absence of the sufficient conditions for point identification and Section 3.3 provides a discussion of further threats to point identification.

3.1 Definition of point identification

So as to define point identification, it is necessary to define observational equivalence of strategic behavior types. If there are two strategic behavior types that *are not* observationally equivalent, then at least, in principle, in at least some games, those two strategic behavior types could generate different observed behavior, and, therefore, be distinguished from each other. Conversely, if there are two strategic behavior types that *are* observationally equivalent, then there are no games in which those two strategic behavior types would generate different observed behavior. Therefore, it is impossible to distinguish between observationally equivalent strategic behavior types.

It follows that any point identification result can *at most* be expected to achieve point identification up to observational equivalence of strategic behavior types. But, by definition, point identification up to observational equivalence is enough to answer any interesting question about behavior, because point identification up to observational equivalence exhausts the relevant information needed to understand the behavior generated by the strategic behavior types.

DEFINITION 1 (Observational Equivalence of Strategic Behavior Types). The quantities $\Theta_1 = (\Lambda_1, \Delta_1, P_1)$ and $\Theta_2 = (\Lambda_2, \Delta_2, P_2)$ are observationally equivalent if

- (i) it holds that $\Lambda_1 = \Lambda_2$,
- (ii) it holds that $\Delta_1 \mathbb{1}[\sum_{k \in \mathcal{M}} \Lambda_1(k) > 0] = \Delta_2 \mathbb{1}[\sum_{k \in \mathcal{M}} \Lambda_2(k) > 0]$,
- (iii) it holds that $P_1 1[\Delta_1 > 0] 1[\sum_{k \in \mathcal{M}} \Lambda_1(k) > 0] = P_2 1[\Delta_2 > 0] 1[\sum_{k \in \mathcal{M}} \Lambda_2(k) > 0].$

By the above definition, two strategic behavior types are observationally equivalent if they use the decision rules with the same probability (i.e., Condition (i)), make computational mistakes with the same probability provided that the types actually use decision rules subject to computational mistakes (i.e., Condition (ii)), and make computational mistakes with the same magnitude provided that the types actually use decision

¹⁶If the model were not point identified, inference following Chernozhukov, Hong, and Tamer (2007), Beresteanu and Molinari (2008), Rosen (2008), Andrews and Soares (2010), Bugni (2010), Canay (2010), Romano and Shaikh (2010), Kline (2011), or Kline and Tamer (2016), among others, would be necessary.

rules subject to computational mistakes and make computational mistakes with positive probability (i.e., Condition (iii)). It is not possible to require that observationally equivalent types have the same probability of making computational mistakes if those types never use decision rules subject to computational mistakes, because in that case the probability of making a computational mistake has no observable implications *in any game*.¹⁷ Similarly, it is not possible to require that observationally equivalent types have the same magnitude of computational mistakes if those types never use decision rules subject to computational mistakes or never make computational mistakes, because in that case the magnitude of computational mistakes has no observable implications *in any game*.

Then the following statement is the definition of point identification.

DEFINITION 2 (Point Identification of Model Parameters). The model parameters are point identified if, for any specifications $\{\Theta_{0r}, \pi_0(r)\}_{r=1}^{\tilde{R}_0}$ and $\{\Theta_{1r}, \pi_1(r)\}_{r=1}^{\tilde{R}_1}$ of the model parameters that satisfy the assumptions and also are such that

(i) both specifications $\{\Theta_{0r}, \pi_0(r)\}_{r=1}^{\tilde{R}_0}$ and $\{\Theta_{1r}, \pi_1(r)\}_{r=1}^{\tilde{R}_1}$ generate the observable data,

(ii) it holds that $\pi_0(\cdot) > 0$ and $\pi_1(\cdot) > 0$,

(iii) the strategic behavior rules Θ_{0r} and $\Theta_{0r'}$ are not observationally equivalent for all $r \neq r'$, and Θ_{1r} and $\Theta_{1r'}$ are not observationally equivalent for all $r \neq r'$,

then $\tilde{R}_0 = \tilde{R} = \tilde{R}_1$ and there is a permutation ϕ of $\{1, 2, ..., \tilde{R}\}$ such that for each $r = 1, 2, ..., \tilde{R}$, it holds that $\pi_0(r) = \pi_1(\phi(r))$ and Θ_{0r} is observationally equivalent to $\Theta_{1\phi(r)}$.

This is the standard definition of point identification, adjusted for two issues. First, point identification can only be up to observationally equivalent strategic behavior types, as discussed above. This concerns parameters relating to computational mistakes, which are assumed known by the econometrician when the model is specified to have no computational mistakes, and otherwise might be viewed as "nuisance parameters." And second, point identification can only be up to permutations of the labeling of the strategic behavior types, because the labeling has no observable implication. As in any model with types, it is not possible to identify which strategic behavior type is truly type r since being type r rather than type r' has no observable implication.

The condition that $\pi(\cdot) > 0$ is required because it is not possible to point identify the strategic behavior types that are used with zero probability. Types that are used by zero percent of the population have no observable implication. So, in a specification that has \tilde{R} strategic behavior types, it is assumed that indeed all \tilde{R} types are used with positive probability. This can be taken as the definition of a specification using \tilde{R} strategic behavior types, ruling out using a type with zero probability. Moreover, the condition that the

¹⁷If observationally equivalent types were required to have the same probability of making computational mistakes even if the types never use decision rules subject to computational mistakes, then two strategic behavior types that generate the same behavior (i.e., two types that use the decision rules with the same probabilities, never use decision rules subject to computational mistakes, and have different probabilities of making a computational mistake) would be defined as *not* observationally equivalent.

strategic behavior types in a specification are not observationally equivalent is required because it is always possible to "split" a strategic behavior type into two identical copies of that type and generate the same observable data, as long as the sum of the probabilities of the use of those two types equals the probability of the use of the original type. By requiring that the types are not observationally equivalent, this uninteresting source of non-identification is ruled out.

3.2 Counterexample to point identification

It is possible to give a counterexample to point identification in the absence of the sufficient conditions established in this paper. This counterexample illustrates the difficulty in distinguishing between across-individual heterogeneity and within-individual heterogeneity.

The counterexample involves two specifications of the parameters. In the first specification, R = 1, and $(\Lambda_1(\text{NE}), \Lambda_1(1_{\text{anch}})) = (\frac{1}{2}, \frac{1}{2})$ and $\Delta_1 = 0$. In the second specification, R = 2, with $\pi(r) = \frac{1}{2}$ and $(\Lambda_r(\text{NE}), \Lambda_r(1_{\text{anch}})) = (1[r = 1], 1[r = 2])$, and $\Delta_r = 0$ for $1 \le r \le 2$. There are a total of three types across these two specifications, and no pairs of types are observationally equivalent according to Definition 1.

In the first specification, all subjects use the same strategic behavior rule, and that rule uses the Nash equilibrium and one step of anchored strategic reasoning with equal probability. In the second specification, there are two equally probable strategic behavior rules, and each rule uses just one of the decision rules.

These two specifications generate the same data in any one game: an equally weighted mixture of point masses at the actions associated with Nash equilibrium and one step of anchored strategic reasoning. Consequently, these two specifications cannot be distinguished on the basis of observing subjects play just one game, and, therefore, the parameters of the model are not point identified if the econometrician observes subjects play just one game. This shows that within-individual heterogeneity cannot be detected in data from just one game. The specification involving within-individual heterogeneity results in the same aggregate distribution of observable data in every game as does the specification not involving within-individual heterogeneity. This is because within-individual heterogeneity is a property of individuals, and, therefore, individuals must be observed to play multiple games so as to identify within-individual heterogeneity. This counterexample is unrelated to the additional complications introduced by computational mistakes or unanchored strategic reasoning, which are discussed in Section 3.3.

Similar counterexamples can be shown in the context of data on more than one game, but less than the number of games established as sufficient for point identification. These counterexamples become notationally cumbersome when the number of games is large but not large enough for point identification, but it is possible to provide another relatively simple counterexample when there are two games. By some abuse of notation, consider the parameterized specification that R = 2, with parameters $\pi(r)$ and $(\Lambda_r(\text{NE}), \Lambda_r(1_{\text{anch}})) = (\Lambda_r, 1 - \Lambda_r)$, and $\Delta_r = 0$ for $1 \le r \le 2$. Note that $\pi(2) = 1 - \pi(1)$. The free parameters are $\pi(1)$, Λ_1 , and Λ_2 . The data when G = 2 can be summarized by

the following four observed probabilities concerning the distribution of subjects' behavior across G = 2 games:

(i) The probability that a subject uses Nash in both games: $P(\text{NE}, \text{NE}) = \Lambda_1^2 \pi(1) + \Lambda_2^2(1 - \pi(1))$.

(ii) The probability that a subject uses Nash and then one step of anchored strategic reasoning: $P(\text{NE}, 1_{\text{anch}}) = \Lambda_1(1 - \Lambda_1)\pi(1) + \Lambda_2(1 - \Lambda_2)(1 - \pi(1))$.

(iii) Equally, due to the assumption that behavior is independent across games, so the order of games does not matter, the probability that a subject uses one step of anchored strategic reasoning and then Nash: $P(1_{\text{anch}}, \text{NE}) = \Lambda_1(1 - \Lambda_1)\pi(1) + \Lambda_2(1 - \Lambda_2) \times (1 - \pi(1))$.

(iv) The probability that a subject uses one step of anchored strategic reasoning in both games: $P(1_{\text{anch}}, 1_{\text{anch}}) = (1 - \Lambda_1)^2 \pi(1) + (1 - \Lambda_2)^2 (1 - \pi(1)).$

Consequently, if there are two distinct specifications of $\pi(1)$, Λ_1 , and Λ_2 that give rise to the same numerical values for these four probabilities (i.e., P(NE, NE), $P(\text{NE}, 1_{\text{anch}})$, $P(1_{\text{anch}}, \text{NE})$, and $P(1_{\text{anch}}, 1_{\text{anch}})$), then the model is not point identified. It is a fairly straightforward computational exercise to establish. For just one example, the specification ($\pi(1) = 0.16$, $\Lambda_1 = 0.65$, and $\Lambda_2 = 0.4$) generates the same values for these four probabilities as does ($\pi(1) = 0.3$, $\Lambda_1 = 0.3$, and $\Lambda_2 = 0.5$).

3.3 Further threats to point identification

Section 3.2 is an example of one threat to point identification. There are many other threats.

First, because of computational mistakes, even if a subject does not use the action associated with a particular decision rule, that subject nevertheless may have intended to use that decision rule. For example, a subject might have intended to use Nash equilibrium, but only use an action approximately equal to the Nash equilibrium action, due to a computational mistake. Therefore, it is not enough to check whether a subject uses the associated action so as to check whether that subject used that decision rule.

Second, when multiple decision rules predict the same action in a given game, then based on observing a subject take that action, it is impossible to uniquely determine the decision rule. In particular, any action that is predicted by s' steps of unanchored strategic reasoning is also predicted by s steps of unanchored strategic reasoning for $0 \le s \le s'$, as discussed in Section 2.2.2.

Third, the observed actions are not necessarily identically distributed across games. For example, it could be that in one game, a particular range of actions is consistent with both zero and one steps of unanchored strategic reasoning, but in another game, that same range of actions is consistent with only zero steps of unanchored strategic reasoning. Consequently, the probability of observing actions in that range would be different across the two games, even holding fixed the probabilities that subjects use the various decision rules. Therefore, observed actions across games are not necessarily identically distributed, despite the fact that the use of decision rules is identically distributed across games per $\lambda_i(\cdot)$.

4. Sufficient conditions for point identification of all model parameters

This section provides the main sufficient conditions for point identification of all unknown model parameters, in the sense of Definition 2. Because the main sufficient conditions for point identification concern the properties of the games that subjects are observed to play, the identification result can be interpreted as a result on *experimental design*. An econometrician with the goal of identifying the decision rules should conduct an experiment that has subjects play games that satisfy the conditions of the identification result. Mechanically, estimation is straightforward under the sufficient conditions for point identification, and it proceeds by maximizing the likelihood derived in Appendix A.

The sufficient conditions for point identification must be at least as strong as any necessary condition for point identification. It is a necessary condition for point identification that each pair of decision rules in the model makes distinct predictions relative to the games in the experiment. Otherwise, if two decision rules in the model make the same predictions in all of the games in the experiment, then obviously those two decision rules are observationally equivalent relative to the games in the experiment. Also, per the counterexample in Section 3.2, a necessary condition for point identification is that each subject is observed to play multiple games.

To distinguish between the use of different numbers of steps of unanchored strategic reasoning, it is necessary that the different numbers of steps of unanchored strategic reasoning make distinct predictions in at least some of the games in the experiment. Section 2.2.2 discussed the fact that, in every game, some actions are consistent with multiple different numbers of steps of unanchored strategic reasoning. Nevertheless, it is possible to distinguish between the use of different numbers of steps of unanchored strategic reasoning, because some actions are *inconsistent* with certain numbers of steps of unanchored strategic reasoning.

Define the set $U_{jg}(s, \varepsilon)$ to be a (possibly empty) set of actions for player *j* in game *g* that are consistent with *s* steps of unanchored strategic reasoning, are not consistent with $s' \in \mathcal{U}$ with s' > s steps of unanchored strategic reasoning, and collectively will be taken with zero probability by subjects who use any decision rule $k \in \mathcal{M}$ and possibly make a computational mistake of magnitude at most ε . The set $U_{jg}(s, \varepsilon)$ can be written as

$$U_{jg}(s,\varepsilon) = \begin{cases} U_{jg}^1(s,\varepsilon) & \text{if } \Sigma_{jg}^s \text{ is not a finite set,} \\ U_{jg}^0(s) & \text{if } \Sigma_{jg}^s \text{ is a finite set,} \end{cases}$$

where

$$\begin{split} U_{jg}^{1}(s,\varepsilon) &= \Sigma_{jg}^{s} \cap \bigcap_{k \in \mathcal{M}} \left[c_{jg}(k) - \varepsilon \left(\alpha_{Ug}(j) - \alpha_{Lg}(j) \right), c_{jg}(k) + \varepsilon \left(\alpha_{Ug}(j) - \alpha_{Lg}(j) \right) \right]^{C} \\ &\cap \bigcap_{s' > s, s' \in \mathcal{U}} \left(\Sigma_{jg}^{s'} \right)^{C}, \\ U_{jg}^{0}(s) &= \Sigma_{jg}^{s} \cap \bigcap_{k \in \mathcal{M}} \left\{ c_{jg}(k) \right\}^{C} \cap \bigcap_{s' > s, s' \in \mathcal{U}} \left(\Sigma_{jg}^{s'} \right)^{C}. \end{split}$$



FIGURE 1. Stylized graphical depiction of Assumption 4.1. This figure complements the discussion of Assumption 4.1, showing a stylized depiction of the arrangement of various quantities in the action space. In this depiction, subjects might use zero or one steps of unanchored strategic reasoning, or one step of anchored strategic reasoning, or Nash equilibrium. Recall that zero steps of anchored strategic reasoning is the same as zero steps of unanchored strategic reasoning, at least under the uniform selection rule on unanchored strategic reasoning.

Let $R_{jg}(s, s', \varepsilon)$ be the probability of $U_{jg}(s, \varepsilon)$ under the distribution with density $\zeta_{jg}^{s'}(\cdot)$ with respect to the appropriate dominating measure on $\Sigma_{jg}^{s'}$ from Section 2.2.2.¹⁸ By construction, $R_{jg}(s, s', \varepsilon) = 0$ if s' > s and $s' \in \mathcal{U}$. Let $U_{jg}(s) = U_{jg}(s, \overline{\rho})$, where $\overline{\rho}$ comes from Assumption 2.2. Also, let $\Omega_{jg} = \alpha_{Ug}(j) - \alpha_{Lg}(j)$.

The addition of Assumption 4.1 is sufficient for point identification. A stylized depiction of the assumption is provided in Figure 1, which shows the arrangement of various quantities in the action space in the case that $\Sigma_{1g}^s = [c_{L1g}(s), c_{U1g}(s)]$. Recall from Section 2.7 that, without loss of generality, the subjects in the data set are always player 1 in the games. Assumption 4.1 is discussed in more detail in the context of the empirical application in Section 5.

ASSUMPTION 4.1 (Conditions on the Games). The data set includes at least 2R - 1 games, such that each game g of those 2R - 1 games satisfies all of the following conditions:

- (i) It holds that $\Omega_{1g} > 0$.
- (ii) For each $k \in \mathcal{M}$ and $k' \in \mathcal{M}$ such that $k \neq k'$, $|c_{1g}(k) c_{1g}(k')| > 2\overline{\rho}\Omega_{1g}$.
- (iii) For each $k \in \mathcal{M}$ and $s \in \mathcal{U}$ such that Σ_{1g}^s is a finite set, $c_{1g}(k) \notin \Sigma_{1g}^s$.
- (iv) For each $k \in \mathcal{M}$, $\overline{\rho}\Omega_{1g} < \max\{\alpha_{Ug}(1) c_{1g}(k), c_{1g}(k) \alpha_{Lg}(1)\}$.
- (v) For each $s \in \mathcal{U}$, $R_{1g}(s, s, \overline{\rho}) > 0$.

Note that Assumption 4.1 requires the data set to include at least 2R - 1 games that *simultaneously* satisfy each of the conditions, which is stronger than the condition that,

¹⁸For example, under the uniform selection rule, if $s' \leq s$ and $\Sigma_{jg}^{s'} = [c_{Ljg}(s'), c_{Ujg}(s')]$ is a nondegenerate interval, then $R_{jg}(s, s', \varepsilon)$ is the ratio of the Lebesgue measure of $U_{jg}(s, \varepsilon)$ to $c_{Ujg}(s') - c_{Ljg}(s')$.

for each of the conditions, the data set includes at least 2R - 1 games that satisfy that condition.

Condition (i) requires that the game has nondegenerate action space. If the game had a degenerate action space, then all solution concepts and decision rules would make the same prediction and, therefore, would be observationally equivalent.

Condition (ii) requires that the game be such that the actions predicted by decision rules subject to computational mistakes are far enough apart from each other, relative to the largest possible computational mistakes, so that a subject who uses decision rule $k \in \mathcal{M}$ will take a different action than a subject who uses decision rule $k' \in \mathcal{M}$ for $k' \neq k$, even if the subjects make computational mistakes. Note that if the econometrician specifies the model to have no computational mistakes (i.e., $\overline{\rho} = 0$), this requires simply that $c_{1g}(k) \neq c_{1g}(k')$. Despite this condition, note that it is not necessarily possible to determine the intended decision rule of a subject even if a subject is observed to take an action close to an action predicted by a particular decision rule $k^* \in \mathcal{M}$, because it is still possible that the subject used some number of steps of unanchored strategic reasoning that resulted in taking an action close to the action predicted by decision rule k^* . Moreover, it is not possible to determine the probability that a subject intends to use a decision rule $k^* \in \mathcal{M}$ by checking how often the subject takes the action exactly predicted by decision rule k^* , because with unknown probability the subject will make a computational mistake. In Figure 1, this condition is reflected by the fact that $[c_{1g}(NE) - \overline{\rho}\Omega_{1g}, c_{1g}(NE) + \overline{\rho}\Omega_{1g}]$ is disjoint from $[c_{1g}(1_{anch}) - \overline{\rho}\Omega_{1g}, c_{1g}(1_{anch}) + \overline{\rho}\Omega_{1g}]$.

Condition (iii) requires that the game be such that if it happens that *s* steps of unanchored strategic reasoning predict a finite set of actions, then the actions predicted by decision rules subject to computational mistakes are not equal to one of the finitely many actions predicted by *s* steps of unanchored strategic reasoning. In particular, this is used to distinguish between anchored and unanchored strategic reasoning, because it implies that the actions predicted by decision rules subject to computational mistakes will not arise with positive probability due to the use of unanchored strategic reasoning. In Figure 1, this condition is not relevant as it is assumed that Σ_{1g}^{s} is a nondegenerate interval.

Condition (iv) requires that the game be such that the actions predicted by decision rules subject to computational mistakes are sufficiently far from at least one of the boundaries of the action space. As a consequence, there will be some actions between the largest (or, respectively, smallest) action that arises due to computational mistakes and the upper bound (or, respectively, lower bound) of the action space. Otherwise, it would not be possible to determine the true magnitude of computational mistakes. It allows that the action predicted by a decision rule subject to computational mistakes equals one of the boundaries of the action space. In Figure 1, this condition is reflected by the fact that $[c_{1g}(NE) - \overline{\rho}\Omega_{1g}, c_{1g}(NE) + \overline{\rho}\Omega_{1g}]$ and $[c_{1g}(1_{anch}) - \overline{\rho}\Omega_{1g}, c_{1g}(1_{anch}) + \overline{\rho}\Omega_{1g}]$ are strictly contained in the action space.

Condition (v) requires the game to be such that for each number of steps of unanchored strategic reasoning $s \in U$, there is a set of actions that can only arise from s or fewer steps of unanchored strategic reasoning. This helps to identify the probability of using s + 1 steps of unanchored strategic reasoning, by the difference between the probabilities of using s or fewer and using s + 1 or fewer steps of unanchored strategic reasoning. In Figure 1, this condition is illustrated by $U_{1g}(0)$, which can arise from the use of zero but not one step of unanchored strategic reasoning, and also not the use of other decision rules.

Assumption 4.1 requires that the econometrician observe each of the subjects play at least 2R - 1 games satisfying these conditions. This is necessary to avoid the threat to point identification that was described in Section 3.2.

The econometrician must also observe subjects play at least one game that satisfies some of the above conditions, and a condition described in the following assumption.

ASSUMPTION 4.2 (Conditions on at Least One Game). The econometrician observes in the data set at least one game g satisfying Conditions (i), (ii), and (iv) in Assumption 4.1, and the extra condition that, for each $k \in M$ and $s \in U \cup \{0_{unanch}\}$, one of the following conditions holds:

- (i) $[c_{1g}(k) \overline{\rho}\Omega_{1g}, c_{1g}(k) + \overline{\rho}\Omega_{1g}]$ is a subset of Σ_{1g}^s ,
- (ii) $[c_{1g}(k) \overline{\rho}\Omega_{1g}, c_{1g}(k) + \overline{\rho}\Omega_{1g}]$ is disjoint from Σ_{1g}^s ,
- (iii) $[c_{1g}(k), c_{1g}(k) + \overline{\rho}\Omega_{1g}]$ is a subset of Σ_{1g}^s and $c_{1g}(k) = \alpha_{Lg}(1)$,
- (iv) $[c_{1g}(k) \overline{\rho}\Omega_{1g}, c_{1g}(k)]$ is a subset of Σ_{1g}^s and $c_{1g}(k) = \alpha_{Ug}(1)$.

Conditions (i)–(iv) requires that the range of possible computational mistakes from any decision rule $k \in \mathcal{M}$ cannot overlap the boundary of the range of predictions from any number of steps of unanchored strategic reasoning. This assumption is used to identify the magnitude of computational mistakes by inspecting whether actions slightly closer to the actions predicted by decision rules subject to computational mistakes are more likely than those slightly further away. This assumption guarantees that over the relevant range of possible computational mistakes, the use of unanchored strategic reasoning cannot either mimic or, alternatively, mask computational mistakes. Parts (i) and (ii) of Assumption 4.2 can be viewed, roughly, as meaning that Assumption 4.2 is satisfied whenever the actions associated with the strategies in \mathcal{M} are suitably distinct from the boundaries of the sets of actions associated with unanchored strategic reasoning. Parts (iii) and (iv) of Assumption 4.2 allows that an action associated with a strategy in \mathcal{M} is on the boundary of the action space. Recall from above that $\overline{\rho} = 0$ whenever computational mistakes are ruled out. In that case, note that logically either (i) or (ii) must be true, since the singleton $c_{1g}(k)$ must either be a subset or disjoint from any given set. In Figure 1, this is reflected by the fact that $c_{1g}(NE)$ and $c_{1g}(1_{anch})$ are distinct from the boundaries of the sets of actions associated with unanchored strategic reasoning, hence $[c_{1g}(NE) - \overline{\rho}\Omega_{1g}, c_{1g}(NE) + \overline{\rho}\Omega_{1g}]$ and $[c_{1g}(1_{anch}) - \overline{\rho}\Omega_{1g}, c_{1g}(1_{anch}) + \overline{\rho}\Omega_{1g}]$ are contained in both Σ_{1g}^0 and Σ_{1g}^1 . As with the other assumptions, this assumption is further discussed in the context of the empirical application in Appendix E.

The following theorem establishes that the model is point identified under the above assumptions. The lengthy proof of this theorem is collected in Appendix D. A stylized sketch of the proof is provided in Section 4.1.

THEOREM 4.1. Under Assumptions 2.1, 2.2, 4.1, and 4.2, the parameters of the model are point identified in the sense of Definition 2.

This theorem does not imply that only the games that satisfy the conditions in Assumptions 4.1 or 4.2 are informative about model parameters or that only such games should be used in estimation. All games should be used in estimation for the purposes of maximizing the efficiency of the estimator relative to the available data. Theorem B.1 in Appendix B establishes sufficient conditions for point identification of all unknown parameters except for those related to the *magnitude* of computational mistakes, under weaker conditions than used by Theorem 4.1.

The assumptions do not impose the existence of within-individual heterogeneity, or the existence of across-individual heterogeneity, or the use of any specific decision rule from those described in Section 2.2. In other words, the assumptions do not impose that any individual use the Nash equilibrium or that any individual use unanchored strategic reasoning, and so forth. Therefore, the identification result automatically applies to special cases of the model involving some but not all of those features. As a consequence, estimation of the model does not impose the existence of those features. Rather, estimates of the model can be used to test for the existence of those features. In particular, the identification result can be used to identify a model involving all such features except within-individual heterogeneity if the econometrician knows that each subject can be characterized by the use of just one decision rule. Sections 3.2 and 3.3 show that these features of the model are independent complications of the identification problem.

4.1 Sketch of proof

The formal proof is lengthy and technical, but it is possible to provide a sketch of the proof. The discussion of Assumptions 4.1 and 4.2 already describes the sources of identification, and this sketch describes how they are formalized in the proof. This sketch states without justification the main claims that are nontrivial to prove, and proving those claims comprises a significant portion of the proof.

It can be shown that a vector of probabilities of events related to the observed actions (e.g., the probability of an observed action within a certain range) in game g due to a subject who uses strategic behavior rule θ , $P_{g,\theta}$, can be written as a matrix Q_g that depends on the structure of game g times a vector that is a known function $\eta^*(\cdot)$ (defined in Appendix D) of strategic behavior rule θ . So $P_{g,\theta} = Q_g \eta^*(\theta)$. $P_{g,\theta}$ is not observable, since the population uses more than one strategic behavior rule. Critically, Q_g is nonsingular under the identification assumptions, although that is not obvious and requires a lengthy proof. That implies that if it *were* possible to observe $P_{g,\theta}$, then it would be possible to recover $\eta^*(\theta)$. Let \mathcal{G} be a subset of games of $\{1, 2, \ldots, G\}$. Let $\mathcal{G}(p)$ be the pth smallest element of \mathcal{G} and let $\mathcal{G}_p = \{\mathcal{G}(1), \ldots, \mathcal{G}(p)\}$.

Then, by the algebra of the Kronecker product, the joint distribution of those events across games in the first *p* games out of *G* is $P_{G,\theta,p} \equiv \bigotimes_{g \in \mathcal{G}_p} P_{g,\theta} = \bigotimes_{g \in \mathcal{G}_p} (Q_g \eta^*(\theta)) = (\bigotimes_{g \in \mathcal{G}_p} Q_g)(\bigotimes^p \eta^*(\theta)) = Q_G^{(p)} \eta^*(\theta)^{(p)}$. Again by the algebra of the Kronecker product, $Q_G^{(p)} \equiv \bigotimes_{g \in \mathcal{G}_p} Q_g$ is nonsingular since each Q_g is nonsingular. Let $P_{G,\theta} = (1, P_{G,\theta,1}, \dots, Q_g^{(p)}) = (1, P_{G,\theta,1}, \dots, Q_g^{(p)})$.

 $P_{\mathcal{G},\theta,|\mathcal{G}|}$). Let $\eta^*(\theta)^{(0)} = 1$ and let $\eta^*(\theta)^{(p)} = \eta^*(\theta) \otimes \cdots \otimes \eta^*(\theta)$ be the *p*-times Kronecker product. Let $\overline{\eta}^*(\theta) = (1, \eta^*(\theta)^{(1)}, \dots, \eta^*(\theta)^{(|\mathcal{G}|)})$. Let $Q_{\mathcal{G}}$ be the block diagonal matrix with blocks along the diagonal equal to $Q_{\mathcal{G}}^{(0)}, \dots, Q_{\mathcal{G}}^{(|\mathcal{G}|)}$, which is nonsingular because each term is nonsingular, and let $P_{\mathcal{G},\theta} = Q_{\mathcal{G}}\overline{\eta}^*(\theta)$.

Suppose that the true parameters of the data generating process are rules $\Theta_{0,1}, \ldots, \Theta_{0,R}$, which are used by $\pi_0(1), \ldots, \pi_0(R)$ percent of the population. Let Y_0^* be a matrix that stacks $(\overline{\eta}^*(\Theta_{0,r}))$ for $r = 1, 2, \ldots, R$ as its columns. So then the observable joint distribution of those events across games is $P_{\mathcal{G}} = Q_{\mathcal{G}}Y_0^*\pi_0$. Suppose that another specification of the parameters with rules $\Theta_{1,1}, \ldots, \Theta_{1,R}$, which are used by $\pi_1(1), \ldots, \pi_1(R)$ percent of the population, is observationally equivalent, so that there is an Y_1^* derived from those parameters so that $P_{\mathcal{G}} = Q_{\mathcal{G}}Y_1^*\pi_1$. Then it would hold that $0 = Q_{\mathcal{G}}\overline{Y}^*\overline{\pi}$, where \overline{Y}^* collects the unique columns of Y_0^* and Y_1^* . Correspondingly, $\overline{\pi}$ collects the difference between π_0 and π_1 . The value of π_0 (or π_1) for a strategic behavior rule that does not appear in specification 0 (or 1) is by convention zero, reflecting the fact that that strategic behavior rule is used by zero percent of the population under specification 0 (or 1).

Therefore, $\overline{\pi}$ is in the null space of $Q_{\mathcal{G}}\overline{Y}^*$. It can be shown as a nontrivial claim under the conditions of the identification results that \overline{Y}^* has full column rank. This step critically uses the fact that the econometrician observes at least 2R - 1 games satisfying the conditions of Assumption 4.1. Since $Q_{\mathcal{G}}$ is nonsingular, it follows that $Q_{\mathcal{G}}\overline{Y}^*$ has full column rank, so it must be that $\overline{\pi} = 0$, so that the columns of Y_0^* and Y_1^* are the same up to permutations of the order of the columns. That implies, up to permutations of the labels, that $\eta^*(\cdot)$ applied to the strategic behavior rules in specification 0 is the same as $\eta^*(\cdot)$ applied to the issues relating to possible lack of observable implications of parameters relating to computational mistakes accounted for in Definition 1. So the parameters are point identified in the sense of Definition 2.

5. Empirical application

The empirical application shows that the features of the model are empirically relevant in the context of a well known and representative experimental design, motivating the main contributions of the paper: proposing and establishing identification of the model. Specifically, the empirical application establishes evidence for within-individual heterogeneity and unanchored strategic reasoning.

5.1 Data

The data for the empirical application comes from the two-player guessing game experiment conducted in Costa-Gomes and Crawford (2006). The following text briefly describes the data. The data concern N = 88 subjects, each of whom plays G = 16 games. An important feature of the experimental design is that the subjects face new opponents in each game and do not learn the actions of their opponents until after the conclusion of the experiment. This eliminates basically any role for learning or specializing their play against their perception of their current opponent. This is consistent with the broader view that non-equilibrium models are best studied in a setting without learning.¹⁹ The empirical analysis of these data is quite different in Costa-Gomes and Crawford (2006), because of the difference in models. In Costa-Gomes and Crawford (2006), as representative of the literature, each subject is assumed to have no within-individual heterogeneity, and the model does not include unanchored strategic reasoning, which means that the main result of estimating the model is essentially assigning each subject to its level out of the level-k model.²⁰ The analysis in this current paper does not use the novel information search data that are also studied in Costa-Gomes and Crawford (2006), simply because the data set without the information search data is more representative of the literature, since most studies do not (yet) use such data. Because of these fundamental differences, the analysis in this current paper is not in any sense an attempt to "replicate" the results of Costa-Gomes and Crawford (2006), though Section 5.4 does show how the results are related. Rather, the analysis is intended to show the empirical relevance of the theoretical results of this paper (proposing and point identifying the model) in the context of a well known and representative experimental design.²¹

All of the games are two-player guessing games, which are related to the beauty contests studied by Nagel (1995), Ho, Camerer, and Weigelt (1998), and Bosch-Domenech, Montalvo, Nagel, and Satorra (2002), among others, simply in the sense that all involve the need to guess what the opponents will guess. In a two-player guessing game, two players simultaneously make a guess. The utility function for a player *j* in game *g* is a decreasing function of the difference between its own guess (a_j) , and that player's target (p_{jg}) times the guess of the other player (a_{-j}) . In game *g*, the action space for player *j* is $[\alpha_{Lg}(j), \alpha_{Ug}(j)]$. The utility function for player *j* in game *g* is

$$u_{jg}(a_1, a_2) = \max\{0, 200 - |a_j - p_{jg}a_{-j}|\} + \max\{0, 100 - \frac{|a_j - p_{jg}a_{-j}|}{10}\}$$

For example, if a player's target is $\frac{2}{3}$, then that player's utility is maximized, holding fixed the other player's guess, by guessing two-thirds of the other player's guess. As dis-

¹⁹Another important feature of the experimental design is that the experiment involves only 8 different two-player games in the traditional sense of the definition of a game. However, each subject plays each game once in each of the player roles (i.e., row player and column player), so that each subject plays 16 times. Each such game times player role pair is denoted a separate game. Essentially the same convention is maintained in Costa-Gomes and Crawford (2006).

²⁰The model in Costa-Gomes and Crawford (2006) also allows Nash equilibrium, and certain "dominance" or "sophisticated" strategies (which are rare). Note that the dominance type is distinct from unanchored strategic reasoning despite the fact that unanchored strategic reasoning relates to iterated dominance. Specifically, the definition is such that all of the dominance or sophisticated types make a unique prediction, more similar to the unique predictions made by anchored strategic reasoning and Nash equilibrium, but fundamentally unlike unanchored strategic reasoning. See Costa-Gomes and Crawford (2006, Table 5) for the specific unique actions predicted by the dominance and sophisticated strategies. Costa-Gomes and Crawford (2006) also check for consistency with iterated deletion of dominated strategies, in the sense discussed in Section 2.2.2. The model in Costa-Gomes and Crawford (2006) also allows computational mistakes, somewhat similarly to the treatment of computation mistakes in the model in this current paper.

²¹Using prior experimental data also avoids the time and financial cost of running an experiment that would, in any case, attempt to be representative of other experiments. So since the point is not to innovate the experimental design, it seems to make most sense to use prior experimental data.

Game Specification					Predictions of Decision Rules					
	Player 1		Player 2		Targets		Anchored Reasoning		Unanchored Reasoning	
g	$\alpha_L(1)$	$\alpha_U(1)$	$\alpha_L(2)$	$\alpha_U(2)$	p_1	p_2	$c_1(1_{\text{anch}})$	$c_1(2_{\text{anch}})$	Δ_1^1	$c_1(NE)$
1	100	500	100	900	0.70	0.50	350	105	[100, 500]	100
2	100	900	100	500	0.50	0.70	150	175	[100, 250]	100
3	100	900	300	500	0.50	0.70	200	175	[150, 250]	150
4	300	500	100	900	0.70	0.50	350	300	[300, 500]	300
5	300	500	300	900	1.50	1.30	500	500	[450, 500]	500
6	300	900	300	500	1.30	1.50	520	650	[390, 650]	650
7	300	900	300	900	1.30	1.30	780	900	[390, 900]	900
8	300	900	300	900	1.30	1.30	780	900	[390, 900]	900
9	100	900	100	500	0.50	1.50	150	250	[100, 250]	100
10	100	500	100	900	1.50	0.50	500	225	[150, 500]	150
11	300	900	100	900	0.70	1.30	350	546	[300, 630]	300
12	100	900	300	900	1.30	0.70	780	455	[390, 900]	390
13	300	500	100	900	0.70	1.50	350	420	[300, 500]	500
14	100	900	300	500	1.50	0.70	600	525	[450, 750]	750
15	100	500	100	500	0.70	1.50	210	315	[100, 350]	350
16	100	500	100	500	1.50	0.70	450	315	[150, 500]	500

TABLE 1. Experimental design.

Note: Some numbers are rounded to the nearest integer in this table to avoid clutter. However, in the econometric analysis, the unrounded numbers are used. The numerical values for these strategies are derived using the method described in the text.

played in Table 1, the 16 games differ along two dimensions: the action spaces and the targets. The experimental design and arrangement of the data set are such that when a subject is observed to play some game *g*, that subject is player 1 in the game.

The strategies corresponding to the various decision rules described in Section 2 for player 1 are in the last columns of the table. Strategies for player 2 are not explicitly shown, but the experimental design described in footnote 19 implies that the strategies of player 2 in even (odd) numbered games are the strategies of player 1 in the previous (next) game in the table. In these games, the Nash equilibrium is indistinguishable from rationalizability, since they imply the same guess (i.e., same pure strategy).

As detailed in Costa-Gomes and Crawford (2006), the derivation of the guesses predicted by anchored strategic reasoning (the level-k model) in these games is straightforward. For example, one step of anchored strategic reasoning amounts to using the best response to the opponent using the uniform distribution over its action space. In these games, that best response is a unique action, and can be derived using the properties of the utility function by noting that the best response given the opponent uses the uniform distribution is the same as the best response given the opponent uses the action at the midpoint of its action space.²² And two steps of anchored strategic reasoning amounts to using the best response to an opponent using the action consistent with one step of anchored strategic reasoning, and so on. Similarly, the derivation of the ranges of

²²So, for example, $c_1(1_{anch}) = 350$ in game g = 1 because the midpoint of player 2's action space is 500 and the target for player 1 is 0.7, so the best response of player 1 is to take action $0.7 \times 500 = 350$.

guesses predicted by unanchored strategic reasoning is also straightforward. Let

$$\chi_{jg}(a) = \begin{cases} \alpha_{Lg}(j) & \text{if } a < \alpha_{Lg}(j), \\ a & \text{if } \alpha_{Lg}(j) \le a \le \alpha_{Ug}(j), \\ \alpha_{Ug}(j) & \text{if } a > \alpha_{Ug}(j). \end{cases}$$

The result is that $\sum_{jg}^{s} = [c_{Ljg}(s), c_{Ujg}(s)]$ is an interval. The biggest guess that player j in game g who uses one step of unanchored strategic reasoning can justify making is $c_{Ujg}(1) = \chi_{jg}(p_{jg}\alpha_{Ug}(-j))$. That is because the biggest justifiable guess is the biggest possible guess of the opponent times the target. If that would be outside the action space, then the boundary of the action space is the biggest justifiable guess. Similarly, the smallest guess that player j in game g who uses one step of unanchored strategic reasoning can justify making is $c_{Ljg}(1) = \chi_{jg}(p_{jg}\alpha_{Lg}(-j))$. More generally, the biggest (respectively, smallest) guess that player j in game g who uses s steps of unanchored strategic reasoning can make is $c_{Ujg}(s) = \chi_{jg}(p_{jg}c_{U,-j,g}(s-1))$ (respectively, $c_{Ljg}(s) = \chi_{jg}(p_{jg}c_{L,-j,g}(s-1))$).

5.2 Nonparametric estimates

It is useful to plot the empirical cumulative distribution functions of the observed actions in each of the games. Figure 2 shows this for game 1. The figures for other games are displayed in Appendix F to save space.²³

The actions predicted by 1, 2, and 3 steps of anchored strategic reasoning, and the Nash equilibrium are displayed at the bottom of the figure, along the horizontal axis. The intervals of actions predicted by 1, 2, and 3 steps of unanchored strategic reasoning are displayed via red endpoints at the top of the figure. The interval of actions predicted by 0 steps of unanchored strategic reasoning is necessarily the entire action space.

Figure 2, and the other estimates in Appendix F, shows clear evidence of mass points corresponding to a small number of actions, and otherwise a roughly continuous distribution of actions. In this game, it appears that there are mass points corresponding to using one and two steps of anchored strategic reasoning and the Nash equilibrium, and otherwise a uniform distribution over the action space. The uniform distribution of actions is exactly consistent with the model with a uniform selection rule on unanchored strategic reasoning, as discussed in Section 2.2.2. Thus, in the empirical application, the selection rule on unanchored strategic reason is defined to be the uniform selection rule.²⁴ In this game, 0 and 1 steps of unanchored strategic reasoning make the same predictions about actions, but in other games displayed in Appendix F, the predictions are different.

²³See Appendix D of Costa-Gomes and Crawford (2006) for a different way to display the actions.

²⁴The defining characteristic of a uniformly distributed random variable is a cumulative distribution function with constant slope, which seems essentially to be the case here, after accounting for the mass points. That is, the displayed empirical cumulative distribution function is essentially that of a mixture of point masses and a uniform distribution over the action space. Uniform distributions over the actions consistent with various numbers of steps of unanchored strategic reasoning also appear in the other figures in Appendix F, consistent with Section 2.2.2.



FIGURE 2. Distribution of actions of subjects in game 1.

5.3 Model specification and estimation results

This subsection discusses the final details of model specification and the estimation results. Estimation proceeds by maximizing the likelihood derived in Appendix A. Despite the somewhat complicated likelihood, maximization of the likelihood using Matlab fmincon optimization appears to give adequate computational performance. Appendix E establishes that the sufficient conditions for identification hold in this application. Section 5.3.1 discusses estimation of *R* based on model selection. Sections 5.3.2 and 5.4 discuss the estimation results.

The estimated model does not allow computational mistakes. As a robustness check, the estimation results that do allow computational mistakes are almost identical, as displayed in Appendix G. It is not surprising that the results that allow computational mistakes are almost identical, based on the following argument involving the figures in Section 5.2 and Appendix F. Note that computational mistakes would imply a higher density of actions in the neighborhoods around the actions associated with the decision rules subject to computational mistakes (i.e., the steps of anchored strategic reasoning or Nash equilibrium), compared to the density of actions slightly further away from the actions associated with those same decision rules. In the figures that display the empirical cumulative distribution functions, that would translate to a greater slope of the empirical cumulative distribution functions in those same neighborhoods, compared

to the slope just outside of those neighborhoods. However, there appears to be no such feature in the figures. Note that this argument is agnostic about the exact model of computational mistakes. Although this paper has specified a particular model of computational mistakes, it seems that any reasonable model of computational mistakes would have similar implications. The actions that do not correspond to anchored strategic reasoning or Nash equilibrium appear better explained by unanchored strategic reasoning, not computational mistakes, as the estimation formalizes.

5.3.1 *Model selection* Economic theory does not predict *R*, the number of strategic behavior types. Therefore, *R* is part of the estimation problem. The selection of *R* is based on comparing the likelihood of the models with different *R* adjusted by a measure of model complexity, penalizing models that have more types and, therefore, more parameters. A generic information criterion is $-2\log L_R(\hat{\theta}_R) + h(R, N)$, where L_R is the likelihood function of the data for the model with *R* types, $\hat{\theta}_R$ is the estimate of the parameters of the model with *R* types, and *h* penalizes model complexity as a function of the number of types and sample size. Models with low values of the information criterion are preferred models.

There is not a uniquely "correct" information criterion, so this paper uses two specifications of *h* that are commonly used in the general statistical literature. Suppose that *S* is the total number of decision rules potentially used by the subjects, per Assumption 2.1. Then there are $g_S(R) = R(S) - 1$ free parameters.²⁵

The specification $h(R, N) = g_S(R) \log(N)$ results in the Bayesian information criterion (e.g., Schwarz (1978)). The specification $h(R, N) = 2(g_S(R)) + \frac{2g_S(R)(g_S(R)+1)}{N-g_S(R)-1}$ results in the corrected Akaike information criterion (e.g., Akaike (1974), Sugiura (1978), and Hurvich and Tsai (1989)). See Konishi and Kitagawa (2008) for details on information criteria. Since the information criteria depend on the unknown parameters only through the likelihood, identifiability of the model parameters is irrelevant. Per Theorems 4.1 and B.1, the model will not necessarily be point identified with *R* too large.

The results of model selection are displayed in Table 2, which shows, for each specification of *R*, the values of the Bayesian and Akaike information criteria, and the Δ difference between the information criterion for that *R* and the information criterion for

R	Bayesian	$\Delta_{ m Bayesian}$	Akaike	$\varDelta_{\mathrm{Akaike}}$
1	12,016.81	631.75	12,007.38	662.73
2	11,686.71	301.65	11,666.72	322.07
3	11,484.37	99.31	11,455.44	110.79
4	11,404.61	19.54	11,368.71	24.06
5	11,385.06	0.00	11,344.65	0.00
6	11,386.63	1.57	11,344.79	0.13
7	11,395.20	10.14	11,355.88	11.23

TABLE 2.	Model	selection.
----------	-------	------------

²⁵There are R - 1 free parameters in $\pi(\cdot)$ and S - 1 free parameters per type from $\Lambda_r(\cdot)$. If computational mistakes were allowed, there would be two more free parameters per type.

the specification of *R* with the smallest value of the information criterion. The *R* with a Δ of zero is preferred by the associated information criterion, since it corresponds to the specification of *R* with the smallest information criterion. The results suggest R = 5 and both criteria show overwhelming support for more than one type, since Δ_{Bayesian} and Δ_{Akaike} for the model with R = 1 are extremely large.

5.3.2 *Parameter estimates* The results of estimating the model are displayed in Table 3. Each row of Table 3 corresponds to one of the estimated types. The first five columns (not counting the *r* column) show the probabilities that type uses the various decision rules described in Section 2.2. The sixth column shows the fraction of the population of that type. Also displayed are 95% confidence intervals. The confidence intervals are estimated according to the standard subsampling algorithm, detailed in the notes to Table 3. Types are listed in decreasing order of the fraction of the population that is that type.

The most common type, 44% of the population, primarily uses zero steps of unanchored strategic reasoning (49%), and also uses one step of unanchored strategic reasoning (31%).

The second most common type, 20% of the population, primarily uses one step of anchored strategic reasoning (70%), and also uses zero steps of unanchored strategic reasoning (15%) and one step of unanchored strategic reasoning (11%).

The third most common type, 15% of the population, primarily uses two steps of anchored strategic reasoning (42%), and also uses one step of anchored strategic reasoning (19%) and one step of unanchored strategic reasoning (24%).

	Anchored	Reasoning	Un	ing		
	1	2	0	1	Nash	Probability
r	$\Lambda_r(1_{\text{anch}})$	$\Lambda_r(2_{\rm anch})$	$\Lambda_r(0_{\text{unanch}})$	$\Lambda_r(1_{\text{unanch}})$	$\Lambda_r(\text{NE})$	of Type $\pi(r)$
1	0.10	0.04	0.49	0.31	0.07	0.44
	(0.07, 0.12)	(0.02, 0.06)	(0.38, 0.56)	(0.22, 0.41)	(0.04, 0.10)	(0.37, 0.56)
2	0.70	0.00	0.15	0.11	0.04	0.20
	(0.52, 0.76)	(0.00, 0.00)	(0.10, 0.28)	(0.06, 0.20)	(0.02, 0.06)	(0.14, 0.31)
3	0.19	0.42	0.11	0.24	0.04	0.15
	(0.00, 0.35)	(0.36, 0.77)	(0.00, 0.20)	(0.00, 0.43)	(0.00, 0.06)	(0.09, 0.26)
4	0.06	0.04	0.04	0.40	0.45	0.15
	(0.03, 0.09)	(0.00, 0.06)	(0.00, 0.08)	(0.33, 0.51)	(0.39, 0.58)	(0.06, 0.23)
5	0.08	0.90	0.00	0.02	0.00	0.06
	(0.00, 0.15)	(0.87, 1.00)	(0.00, 0.00)	(0.00, 0.03)	(0.00, 0.00)	(0.00, 0.08)

TABLE 3. Estimates.

Note: The 95% confidence intervals are reported in parentheses, estimated according to the standard subsampling algorithm for maximum likelihood (e.g., Politis, Romano, and Wolf (1999)) by resampling $N_s = \text{floor}(\frac{2}{3}88) = 58$ people from the data set, without replacement. Conventional asymptotic approximations and bootstraps are likely invalid in this model with these data, because many of the estimated probabilities are zero, which suggests a parameter on the boundary problem.

The fourth most common type, 15% of the population, primarily uses the Nash equilibrium (45%), and also uses one step of unanchored strategic reasoning (40%).

Finally, the least common type, 6% of the population, primarily uses two steps of anchored strategic reasoning (90%), and also uses one step of anchored strategic reasoning (8%).

All types involve within-individual heterogeneity, since no type exclusively uses one decision rule. However, the least common type does have relatively little withinindividual heterogeneity. This shows that allowing within-individual heterogeneity is important. The estimated strategic behavior types generally have the sensible feature that they emphasize the use of just one mode of strategic reasoning (anchored or unanchored). Rules 1 and 4 predominantly use unanchored strategic reasoning, while rules 2 and 5 predominantly use anchored strategic reasoning. Rule 3 shows a more even mix of modes of strategic reasoning. This shows that allowing both modes of strategic reasoning is important, and that different subjects use different modes of strategic reasoning. The fact that the estimates are sensible in this way was not imposed by the model or the estimation method.

5.4 Relationship to prior estimates

There is a relationship between the estimates in Table 3 that allow within-individual heterogeneity and the estimates in Costa-Gomes and Crawford (2006) that do not allow within-individual heterogeneity. Based on the standard level-k model and not allowing within-individual heterogeneity, Costa-Gomes and Crawford (2006) observe that roughly half of the subjects can be assigned their type based on type being "apparent from guesses," which means using the action associated with the type in at least 7 out of the 16 games. Another contribution of the model in this paper is including unanchored strategic reasoning, but this is not included in this discussion because Costa-Gomes and Crawford (2006) focus on the level-k model. From the estimates in Table 3, each of the estimated types has an expected fraction of subjects of that type who would use any given decision rule in at least 7 out of 16 games. Weighting these expected fractions by the proportions of the types, thereby integrating over the types, it is possible to compare the fraction of subjects assigned to each decision rule from the Costa-Gomes and Crawford (2006) estimates to the expected fraction of subjects who would use that same decision rule in at least 7 out of 16 games, according to the estimates in Table 3.

This comparison proceeds by separately considering each decision rule. Costa-Gomes and Crawford (2006) find that 20 subjects (22.7%) are the type to use one step of anchored strategic reasoning. Type 2 (20% of the population) uses one step of anchored strategic reasoning with probability 70%. Using the binomial distribution, such subjects will almost surely use one step of anchored strategic reasoning in at least 7 out of 16 games, and, therefore, will appear to be the type that uses one step of anchored strategic reasoning, explaining the concordance between the estimates of 22.7% and 20%. Other types use one step of anchored strategic reasoning so rarely that such subjects are extremely unlikely to use it in 7 out of 16 games, and thus will not appear to be that type. Costa-Gomes and Crawford (2006) also find that 12 subjects (13.6%) are

the type to use two steps of anchored strategic reasoning. Type 5 (6% of the population) uses two steps of anchored strategic reasoning with probability 90%. Such subjects will almost certainly use two steps of anchored strategic reasoning in at least 7 out of 16 games. Moreover, using the binomial distribution, roughly 54% of type 3 subjects (a type comprising 15% of the population) will use two steps of anchored strategic reasoning in at least 7 out of 16 games. Thus, roughly, based on these estimates there will be $6\% + 54\% \times 15\% \approx 14.1\%$ of subjects who will use two steps of anchored strategic reasoning in at least 7 out of 16 games; hence the concordance between the estimates of 13.6% and 14.1%. Finally, Costa-Gomes and Crawford (2006) find that 8 subjects (9.1%) are the type to use Nash equilibrium. Type 4 (15% of the population) uses Nash equilibrium with probability 45%. Using the binomial distribution, approximately 63% of such subjects will use Nash equilibrium in at least 7 out of 16 games; hence the concordance between the concordance between the estimates of subjects will use Nash equilibrium in at least 7 out of 16 games; hence the concordance between the concordance between the subjects will use Nash equilibrium in at least 7 out of 16 games; hence the concordance between the concordance between the estimates of 9.1% and 9.5% = 63% × 15%. Other types use Nash equilibrium so rarely that such subjects are extremely unlikely to use it in 7 out of 16 games.

Costa-Gomes and Crawford (2006) also estimate a model that essentially has the consequence of assigning the subjects to the type that fits (in a maximum likelihood sense) as a best approximation to their underlying within-individual heterogeneity.

6. Conclusion

This paper proposes a structural model of non-equilibrium behavior in games so as to learn about the solution concepts and decision rules that individuals use to determine their actions. The model allows both anchored and unanchored strategic reasoning, and computational mistakes. Also, the model allows both across-individual and withinindividual heterogeneity. The paper proposes the model and provides sufficient conditions for point identification. As discussed in particular in Section 3, these features of the model interact with each other but nevertheless are independent challenges to identification, so the identification result is a contribution even if some but not all of those features are present in a particular application. Because the sufficient conditions concern the structure of the games that subjects are observed to play, the identification result can be interpreted as a result on experimental design, informing the sorts of experiments that should be run to learn about the solution concepts and decision rules that individuals use. Then the model is estimated on data from an experiment involving two-player guessing games. The application both illustrates the empirical relevance of the features of the model and provides empirical results of independent interest. The results indicate both across-individual heterogeneity and within-individual heterogeneity, and both modes of strategic reasoning.

References

Akaike, H. (1974), "A new look at the statistical model identification." *IEEE Transactions on Automatic Control*, 19 (6), 716–723. [174]

An, Y. (2017), "Identification of first-price auctions with non-equilibrium beliefs: A measurement error approach." *Journal of Econometrics*, 200 (2), 326–343. [142]

Andrews, D. W. and G. Soares (2010), "Inference for parameters defined by moment inequalities using generalized moment selection." *Econometrica*, 78 (1), 119–157. [160]

Aradillas-Lopez, A. and A. Rosen (2013), "Inference in ordered response games with complete information." [142]

Aradillas-Lopez, A. and E. Tamer (2008), "The identification power of equilibrium in simple games." *Journal of Business and Economic Statistics*, 26 (3), 261–283. [142, 147]

Aumann, R. J. (2000), "Game theory." In *Collected Papers*, Vol. I, 47–106, MIT Press, Cambridge, MA. [146]

Bajari, P., H. Hong, and S. P. Ryan (2010), "Identification and estimation of a discrete game of complete information." *Econometrica*, 78 (5), 1529–1568. [142]

Bardsley, N., R. Cubitt, G. Loomes, P. Moffatt, C. Starmer, and R. Sugden (2010), *Experimental Economics: Rethinking the Rules*. Princeton University Press, Princeton, NJ. [155]

Beresteanu, A. and F. Molinari (2008), "Asymptotic properties for a class of partially identified models." *Econometrica*, 76 (4), 763–814. [160]

Bernheim, B. D. (1984), "Rationalizable strategic behavior." *Econometrica*, 52 (4), 1007–1028. [142, 147]

Bosch-Domenech, A., J. G. Montalvo, R. Nagel, and A. Satorra (2002), "One, two, (three), infinity, ...: Newspaper and lab beauty-contest experiments." *American Economic Review*, 92 (5), 1687–1701. [170]

Bugni, F. A. (2010), "Bootstrap inference in partially identified models defined by moment inequalities: Coverage of the identified set." *Econometrica*, 78 (2), 735–753. [160]

Camerer, C. (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton, NJ. [141, 143, 144, 154]

Camerer, C. F., T.-H. Ho, and J.-K. Chong (2004), "A cognitive hierarchy model of games." *Quarterly Journal of Economics*, 119 (3), 861–898. [143]

Canay, I. A. (2010), "EL inference for partially identified models: Large deviations optimality and bootstrap validity." *Journal of Econometrics*, 156 (2), 408–425. [160]

Carnap, R. (1953), "Inductive logic and science." *Proceedings of the American Academy of Arts and Sciences*, 80 (3), 189–197. [149]

Chernozhukov, V., H. Hong, and E. Tamer (2007), "Estimation and confidence regions for parameter sets in econometric models." *Econometrica*, 75 (5), 1243–1284. [160]

Costa-Gomes, M., V. P. Crawford, and B. Broseta (2001), "Cognition and behavior in normal-form games: An experimental study." *Econometrica*, 69 (5), 1193–1235. [143]

Costa-Gomes, M. A. and V. P. Crawford (2006), "Cognition and behavior in two-person guessing games: An experimental study." *American Economic Review*, 96 (5), 1737–1768. [141, 143, 144, 151, 169, 170, 171, 172, 176, 177]

Quantitative Economics 9 (2018)

Crawford, V. P., M. A. Costa-Gomes, and N. Iriberri (2012), "Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications." *Journal of Economic Literature*, 51 (1), 5–62. [143, 144]

Crawford, V. P. and N. Iriberri (2007a), "Fatal attraction: Salience, naivete, and sophistication in experimental 'hide-and-seek' games." *American Economic Review*, 97 (5), 1731– 1750. [143]

Crawford, V. P. and N. Iriberri (2007b), "Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions?" *Econometrica*, 75 (6), 1721–1770. [143]

de Paula, Á. (2013), "Econometric analysis of games with multiple equilibria." *Annual Review of Economics*, 5 (1), 107–131. [142]

Dunker, F., S. Hoderlein, and H. Kaido (2013), "Random coefficients in static games of complete information." [142]

Fox, J. T. and N. Lazzati (2017), "A note on identification of discrete choice models for bundles and binary games." *Quantitative Economics*, 8 (3), 1021–1036. [142]

Fudenberg, D. and J. Tirole (1991), *Game Theory*. MIT Press, Cambridge, MA. [142, 147, 148]

Georganas, S., P. J. Healy, and R. A. Weber (2015), "On the persistence of strategic sophistication." *Journal of Economic Theory*, 159, 369–400. [156]

Gillen, B. J. (2010), "Identification and estimation of level-k auctions." [142]

Haile, P. A., A. Hortaçsu, and G. Kosenok (2008), "On the empirical content of quantal response equilibrium." *American Economic Review*, 98 (1), 180–200. [142]

Ho, T.-H., C. Camerer, and K. Weigelt (1998), "Iterated dominance and iterated best response in experimental 'p-beauty contests'." *American Economic Review*, 88 (4), 947–969. [143, 170]

Hurvich, C. M. and C.-L. Tsai (1989), "Regression and time series model selection in small samples." *Biometrika*, 76 (2), 297–307. [174]

Kline, B. (2011), "The Bayesian and frequentist approaches to testing a one-sided hypothesis about a multivariate mean." *Journal of Statistical Planning and Inference*, 141 (9), 3131–3141. [160]

Kline, B. (2015), "Identification of complete information games." *Journal of Econometrics*, 189 (1), 117–131. [142, 147]

Kline, B. (2016), "The empirical content of games with bounded regressors." *Quantitative Economics*, 7 (1), 37–81. [142]

Kline, B. and E. Tamer (2012), "Bounds for best response functions in binary games." *Journal of Econometrics*, 166 (1), 92–105. [142, 147]

Kline, B. and E. Tamer (2016), "Bayesian inference in a class of partially identified models." *Quantitative Economics*, 7 (2), 329–366. [160]

Konishi, S. and G. Kitagawa (2008), *Information Criteria and Statistical Modeling*. Springer, New York. [174]

Machina, M. J. (1985), "Stochastic choice functions generated from deterministic preferences over lotteries." *The Economic Journal*, 95 (379), 575–594. [155]

McFadden, D. (1981), "Econometric models of probabilistic choice." In *Structural Analysis of Discrete Data With Econometric Applications* (C. F. Manski and D. McFadden, eds.), 198–272, MIT Press, Cambridge, MA. [155]

Myerson, R. B. (1991), Game Theory. Harvard University Press, Cambridge, MA. [146]

Nagel, R. (1995), "Unraveling in guessing games: An experimental study." *American Economic Review*, 85 (5), 1313–1326. [143, 170]

Nash, J. F. (1950), "Equilibrium points in n-person games." *Proceedings of the National Academy of Sciences of the United States of America*, 36 (1), 48–49. [141]

Osborne, M. J. and A. Rubinstein (1994), *A Course in Game Theory*. MIT Press, Cambridge, MA. [146]

Pearce, D. G. (1984), "Rationalizable strategic behavior and the problem of perfection." *Econometrica*, 52 (4), 1029–1050. [142, 147]

Politis, D. N., J. P. Romano, and M. Wolf (1999), Subsampling. Springer, New York. [175]

Rieskamp, J. (2008), "The probabilistic nature of preferential choice." *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34 (6), 1446–1465. [155]

Rieskamp, J., J. R. Busemeyer, and B. A. Mellers (2006), "Extending the bounds of rationality: Evidence and theories of preferential choice." *Journal of Economic Literature*, 44 (3), 631–661. [155]

Romano, J. P. and A. M. Shaikh (2010), "Inference for the identified set in partially identified econometric models." *Econometrica*, 78 (1), 169–211. [160]

Rosen, A. M. (2008), "Confidence sets for partially identified parameters that satisfy a finite number of moment inequalities." *Journal of Econometrics*, 146 (1), 107–117. [160]

Schwarz, G. (1978), "Estimating the dimension of a model." *The Annals of Statistics*, 6 (2), 461–464. [174]

Stahl, D. O. and P. W. Wilson (1994), "Experimental evidence on players' models of other players." *Journal of Economic Behavior & Organization*, 25 (3), 309–327. [143]

Stahl, D. O. and P. W. Wilson (1995), "On players' models of other players: Theory and experimental evidence." *Games and Economic Behavior*, 10 (1), 218–254. [143, 156]

Sugiura, N. (1978), "Further analysis of the data by Akaike's information criterion and the finite corrections." *Communications in Statistics—Theory and Methods*, 7 (1), 13–26. [174]

Quantitative Economics 9 (2018)

Tamer, E. (2003), "Incomplete simultaneous discrete response model with multiple equilibria." *Review of Economic Studies*, 70 (1), 147–165. [142]

Tan, T. C.-C. and S. R. da Costa Werlang (1988), "The Bayesian foundations of solution concepts of games." *Journal of Economic Theory*, 45 (2), 370–391. [142, 147, 148, 150]

Co-editor Rosa L. Matzkin handled this manuscript.

Manuscript received 8 December, 2015; final version accepted 18 January, 2017; available online 21 February, 2017.