DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

GENERAL REVEALED PREFERENCE THEORY

Christopher P. Chambers
University of California San Diego

Federico Echenique
California Institute of Technology

Eran Shmaya
Northwestern University

SOCIAL SCIENCE WORKING PAPER 1332

September 2010

General Revealed Preference Theory

Christopher P. Chambers Federico Echenique Eran Shmaya

Abstract

We provide general conditions under which an economic theory has a universal axiomatization: one that leads to testable implications. Roughly speaking, if we obtain a universal axiomatization when we assume that unobservable parameters (such as preferences) are observable, then we can obtain a universal axiomatization purely on observables. The result "explains" classical revealed preference theory, as applied to individual rational choice. We obtain new applications to Nash equilibrium theory and Pareto optimal choice.

JEL classification numbers: A10,D00

Key words: Falsifiability; Axiomatization; Revealed preference; Testable implications of Nash equilibrium; Model Theory

General Revealed Preference Theory *

Christopher P. Chambers Federico Echenique Eran Shmaya

1 Introduction

Revealed preference (RP) theory is the name given to economic theories whose goal is to understand the testable implications of preference-based choice models when preferences are unobservable. The classical model assumes that maximization of some preference relation governs choices from a budget set, and that this preference relation is unobserved. Modern day economics is more permissive in what it means by RP theory, but the idea remains of characterizing the empirical content of a given economic model when some parameters are unobserved.

RP theory presents a special methodological challenge. In all RP exercises, given an economic theory, one postulates that reality behaves *as if* that theory were true. The goal is to understand the testable implications of the theory. The problem, however, with the as-if statement is that it violates a basic tenet of positivist scientific methodology. Specifically, RP theory is an empirical approach that need not necessarily lead to a falsifiable theory.

RP theory, and the *as-if* approach, is usually formulated existentially, as opposed to universally. A universal theory is falsifiable; an existential theory is not. We can explain the distinction using an example of Popper (1959): Theory E states "There is a white swan," while theory U states "All swans are white." Theory E is existential: It is not falsifiable because no matter how many finite data sets of non-white swans we discover, they do not contradict the existence (somewhere) of a white swan. On the other hand, U is a universal theory because it makes a statement about every individual swan. Theory U is falsifiable because the observation of one non-white swan falsifies the theory.

*Chambers and Echenique acknowledge support from the NSF through grant SES-0751980.

We make a distinction between the RP *formulation* of a theory and the RP theory itself. When we speak of the RP formulation of a theory, we mean a *description* of the theory which is existentially quantified, where the existential quantification operates directly on unobservable objects, usually preferences.[1] For example "there is a rational preference relation generating the observed choices." The RP theory itself is a description of the datasets for which there exist preferences (or other unobservable parameters) for which the data are consistent with those preferences.

Even though we do not observe preferences, a possible test of the theory is to check each possible preference, and see if that preference could generate the observed data. In general, there are an infinite number of possible preferences. Therefore, the RP formulation is as problematic as "there is a white swan," because no matter how many unobservable preferences fail to explain the data, we cannot conclude that the data falsify the theory.

While the RP formulation of classical choice theory is existentially quantified, the RP theory of classical choice is, in fact, falsifiable. This is so because there is a *theorem* establishing the equivalence of the theory specified by the RP formulation and a particular universal axiomatization, where the universal quantification operates directly on observable data. Consider the classical theory of the consumer. The RP formulation states that a consumer behaves as if there is a utility function that is maximized by her choices. We know from Samuelson (1938) and Houthakker (1950) that the RP theory has a universal axiomatization: namely the strong axiom of revealed preference. Note that the strong axiom of revealed preference is an axiom (actually a collection of axioms) describing conditions on observable data. Thus, despite being formulated existentially over unobservable parameters, the classical theory of the consumer is falsifiable. In fact, the Samuelson-Houthakker representation is important precisely because it is universally quantified over observable data, and therefore provides a test of the theory. The same is true of more abstract theories of individual choice; the most general universal characterization is probably the one in Richter (1966).

Our main result is to prove that, under very general conditions, the RP formulation of an economic theory always admits a universal representation. As applications, we study theories of collective choice: Nash equilibrium and Pareto optimal choice, for example. The RP formulation of these theories is existential, and to obtain universal

[1]For readers who know model theory, when we speak of existential quantification, we are really talking about *second order* logic. Existential quantification over preferences is an example of existential quantification over predicates.

characterizations has proven very difficult. Our main result implies that these theories always admit a universal representation.

As a corollary we give conditions under which it can be determined that a theory has been falsified. Under some conditions (which are satisfied for Nash equilibrium and Pareto optimal choice, for example) there is an algorithm that, when presented with a falsifying data set will determine that the theory has been falsified. It cannot, however, certify that a data set validates a theory: this is the classical Popperian dichotomy that says that a scientific theory can never be proven correct, only falsified.

Our results follow from a general theorem involving the statements one can make about observable and theoretical terms. The classical revealed preference theory is existential, as we explained above. However, if we imagine that preferences are observable then we can describe the theory using purely universal statements (this description would be an axiomatization, or a test, of the theory, in the fictitious case where preferences are observable). We prove that, whenever that is the case, the universal axiomatization involving both unobservables and observables can be "projected" onto observables. The result is a universal axiomatization that makes statements purely about observables.

Our results build on basic work on model theory and mathematical logic. We suppose an observer formulates a theory about behavior, and that the formulation follows the framework of standard (first order) mathematical logic. Specifically, the theory is in terms of a relational first order language $\mathcal{L}$. For example, in the classical individual choice model, we may suppose that an observer hypothesizes that there exists a pair of relations $\succeq, \succ$, where $\succ$ is the strict part of $\succeq$ and $\succeq$ is a weak order.[2] These relations govern hypothetical choices, and there is another pair of relations R and P which correspond to the actual observed choices; R and P are usually called the revealed preference relations. The observer hypothesizes that whenever $x \, R \, y$, then $x \succeq y$ and whenever $x \, P \, y$, then $x \succ y$; however, the converse implications need not hold.

It is well known that, given a pair R and P, there exists a pair $\succeq$, $\succ$ conforming to the above hypotheses if and only if there are no cycles involving R and at least one occurrence of P. This latter property is called the *strong axiom of revealed preference*. The strong axiom of revealed preference includes only statements involving R and P, hence it relates directly to observables. Moreover, it can be written with a collection of universal sentences; hence it is falsifiable. Note the general exercise: we assume a theory

[2] A weak order is complete and transitive.

can be described using unobservable relations. We want to test, using the observed relations only, whether or not the theory is valid.

In general, our language $\mathcal{L}$ specifies all relations, both observable and unobservable. Our notion of a theory is formally semantic, and consists of a class of structures. We define a theory as a class of $\mathcal{L}$-structures T axiomatized by some collection of universal sentences. To model the idea that some relations are not observed, we consider a language $\mathcal{F}$, where all relation symbols in $\mathcal{F}$ are also symbols of $\mathcal{L}$. The symbols in $\mathcal{F}$ correspond to observable relations. We define a new theory $F(T)$ from our original theory by considering the class of all structures which are elements of T, where the predicates corresponding to $\mathcal{L}\backslash\mathcal{F}$ have been removed.

Our general result is that so long as T is universally axiomatizable, so is $F(T)$; and further, it is axiomatized by the set of all universal consequences of T involving only relations in $\mathcal{F}$. This is exactly what we see in the case of the strong axiom: absence of cycles including P are the only universal implications of our general theory which involve only the symbols R and P.

We call our approach *general revealed preference theory*, as revealed preference theory presumes that there is a preference relation generating choice behavior, but that this preference relation is unobservable. Most revealed preference exercises can be framed as special cases of the result here. However, our theorem is useful in establishing new results. For example, our general result allows there to be multiple unobserved relations. A recent branch of the revealed preference literature focuses on the empirical content of group choice functions in games. The approach is as follows: a set of players is given, and a set of strategies is given. For any nonempty set of strategies, for each agent, it is imagined that a joint choice is observed from the game form derived from those sets. One can ask whether or not the observed choices can be rationalized as Nash equilibria, or as the Pareto optimal joint choices for some set of preferences. Examples of such papers include Sprumont (2000), Xu and Zhou (2007), Galambos (2009), and Lee (2009). There are older studies of the same questions, also about group choice, but using other solution concepts: see Wilson (1970), Plott (1974), and Ledyard (1986).

Here, we work with a generalized notion of equilibrium which incorporates Nash equilibrium, strong Nash equilibrium (Aumann, 1960), and Pareto optimality as special cases. We show that the theory hypothesizing that there exist *strict* preferences rationalizing the observed choices is always universally axiomatizable, and hence falsifiable in principle.

2 Main Results

2.1 Preliminary definitions

We proceed to give some standard definitions from model theory. Readers with at least a minimal exposure to model theory or mathematical logic will want to skip this section.

We first must specify our language $\mathcal{L}$. The language is a primitive and specifies the *syntax*, or the things we can say.

1 Definition. A *language* $\mathcal{L}$ is given by specifying the following:

1. a set of predicate symbols $\mathcal{R}$ and positive integers n_R for each $R \in \mathcal{R}$
2. a set of constant symbols $\mathcal{C}$.

The symbol $R \in \mathcal{R}$ is meant to denote a n_R-ary relation. Note that we focus here on languages without function symbols, so called *relational languages*.

The semantics are specified by concrete mathematical objects, called *structures*. Structures provide the appropriate framework for interpreting our syntax.

2 Definition. An *$\mathcal{L}$-structure* $\mathcal{M}$ is given by the following:

1. a nonempty set M called the *domain* of M
2. a set $R^{\mathcal{M}} \subseteq M^{n_R}$ for each $R \in \mathcal{R}$
3. an element $c^{\mathcal{M}} \in M$ for each $c \in \mathcal{C}$.

When the language $\mathcal{L}$ is understood, we refer to an $\mathcal{L}$-structure simply as a *structure*. The elements $R^{\mathcal{M}}$ and $c^{\mathcal{M}}$ are called *interpretations* of the corresponding symbols in the language $\mathcal{L}$.

3 Definition. Suppose that $\mathcal{M}$ and $\mathcal{N}$ are $\mathcal{L}$-structures with universes M and N respectively. An *$\mathcal{L}$-embedding* $\eta : \mathcal{M} \to \mathcal{N}$ is a one-to-one map $\eta : M \to N$ that preserves the interpretations of all symbols of $\mathcal{L}$: specifically,

1. $R^{\mathcal{M}}(a_1, ..., a_{m_R})$ if and only if $R^{\mathcal{N}}(\eta(a_1), ..., \eta(a_{m_R}))$ for all $R \in \mathcal{R}$ and $a_1, ..., a_{m_R} \in M$
2. $\eta(c^{\mathcal{M}}) = c^{\mathcal{N}}$ for $c \in \mathcal{C}$.

As a notational convention, we write $R^{\mathcal{M}}(a_1, ..., a_{m_R})$ to mean $(a_1, ..., a_{m_R}) \in R^{\mathcal{M}}$.

4 Definition. An *isomorphism* is a bijective $\mathcal{L}$-embedding.

Appendix A gives definitions of sentence, and of the validity of a sentence in a structure. These notions correspond closely to their conventional meaning in English.

2.2 The Model

For a given first order language $\mathcal{G}$, a $\mathcal{G}$-theory is a class of structures for that language which is closed under isomorphism. We say that a theory T is *axiomatized* by a collection of sentences Σ if T consists of exactly of the structures for which each sentence in Σ is valid. Given two theories, T and T', where $T \subseteq T'$, we say that T is *axiomatized* by a collection of sentences Σ *with respect to* T' if T consists of exactly those structures in T' for which each sentence in Σ is valid.

Let $\mathcal{F} = \langle R_1, \ldots, R_N \rangle$ and

$$\mathcal{L} = \langle R_1, \ldots, R_N, Q_1, \ldots, Q_K \rangle,$$

be languages, where all the R_n and Q_k are predicate symbols. Note that $\mathcal{F} \subseteq \mathcal{L}$.

Let T be an $\mathcal{L}$-theory. Define $F(T)$ to be the class of $\mathcal{F}$-structures $(X^*, R_1^*, ..., R_N^*)$ for which there exist relations $Q_1^*, ..., Q_K^*$ such that $(X^*, R_1^*, ..., R_N^*, Q_1^*, ..., Q_K^*) \in T$. That is, $F(T)$ is the *projection* of T onto the language $\mathcal{F}$.

Given a collection of $\mathcal{L}$-sentences Σ, the collection of *$\mathcal{F}$-implications of* Σ is the collection of all logical implications of $\mathcal{L}$ involving only predicates from $\mathcal{F}$.

5 Theorem. *If T has a universal $\mathcal{L}$-axiomatization, then $F(T)$ has a universal $\mathcal{F}$-axiomatization. Moreover, if Σ is a universal axiomatization of T, then the collection of universal $\mathcal{F}$-implications of Σ is an axiomatization of $F(T)$.*

Proof. We first establish that $F(T)$ is universally axiomatizable, using a theorem of Tarski (1954). We then show what this axiomatization should be.

We want to verify the three conditions of Theorem 1.2 in Tarski (1954). Specifically, conditions (i), (ii), and (iii') in Tarski's paper. To this end, we need to show that $F(T)$ is closed under isomorphism and substructure. Lastly, we need to show that for any totally ordered set Θ and indexed collection of models $\mathcal{M}^\theta \in F(T)$ where if $\theta < \theta'$, then $\mathcal{M}^\theta$ is a substructure of $\mathcal{M}^{\theta'}$, there is $\mathcal{M}^* \in F(T)$ for which each $\mathcal{M}^\theta$ is a substructure of $\mathcal{M}^*$.

The first two conditions (closure under substructures and isomorphism) follow because T satisfies those conditions as it is a universal theory. We prove the third condition.

Let Σ be a universal axiomatization of T.

Let $\mathcal{M}^\theta = (X^\theta, R_1^\theta, ..., R_N^\theta)$, $\theta \in \Theta$, be a monotone class of structures of $F(T)$. That is, Θ is totally ordered, and $\mathcal{M}^\theta$ is a substructure of $\mathcal{M}^{\theta'}$ whenever $\theta < \theta'$. Let $\mathcal{M}^* = (X^*, R_1^*, \ldots, R_N^*)$ be defined so that $X^* = \bigcup_\theta X^\theta$ and $R_k^* = \bigcup_\theta R_k^\theta$ for $k = 1, \ldots, N$.

For each θ, let W^θ be the set of lists of relations $(Q_1^*, \ldots, Q_K^*)$ on X^* such that

$$(X^\theta, R_1^\theta, \ldots, R_N^\theta, Q_1^*|_{X^\theta}, \ldots, Q_K^*|_{X^\theta})$$

is a model of Σ. Note that $W^\theta \neq \varnothing$ as $\mathcal{M}^\theta \in F(T)$.

We claim that if $\theta < \theta'$, then $W^{\theta'} \subseteq W^\theta$.

Let $(Q_1^*, \ldots, Q_K^*) \in W^{\theta'}$ and let $\varphi(x_1, \ldots, x_M) \in \Sigma$. Then by assumption, $(X^{\theta'}, R_1^{\theta'} \ldots, R_N^{\theta'}, Q_1^*|_{X^{\theta'}}, \ldots$ is a model of φ, so $\varphi(x_1^*, \ldots, x_M^*)$ is valid for any $\{x_1^*, \ldots, x_M^*\} \subseteq X^{\theta'}$; in particular, it is valid for any $\{x_1^*, \ldots, x_M^*\} \subseteq X^\theta$. Since for all $i = 1, \ldots, N$, $R_i^\theta = R_i^{\theta'} \cap X^\theta$, φ is valid in $(X^\theta, R_1^\theta, \ldots, R_N^\theta, Q_1^*|_{X^\theta}, \ldots, Q_K^*|_{X^\theta})$. As φ was arbitrary, $(Q_1^*, \ldots, Q_K^*) \in W^\theta$.

Note that if Q is a k-ary relation on X^*, it is a subset of X^k. To this end, regard W^θ as a subset of

$$\mathcal{B} = \{0,1\}^{\Pi_1 X^*} \times \cdots \times \{0,1\}^{\Pi_K X^*},$$

where $\Pi_k X^*$ stands for the product $X^* \times \ldots, \times X^*$, as many times as the order of the predicate Q_k. Note that $\mathcal{B}$, endowed with the product topology, is compact.

We claim that W^θ, viewed as a subset of $\mathcal{B}$, is closed. To see this, let $(Q_1^\lambda, \ldots, Q_K^\lambda)$ be a net in W^θ, converging to $(Q_1^*, \ldots, Q_K^*)$. Let $\varphi(x_1, \ldots, x_M)$ be a formula in Σ and $\{x_1^*, \ldots, x_M^*\} \subseteq X^\theta$. Then by definition of product topology convergence, there exists $\overline{\lambda}$ such that if $\overline{\lambda} < \lambda$, then for $k = 1, \ldots, K$, $Q_k^\lambda(x_1^* \ldots, x_M^*)$ if and only if $Q_k^*(x_1^* \ldots, x_M^*)$. Then, since

$$w = \langle X^\theta, R_1^\theta, \ldots, R_N^\theta, Q_1^\lambda|_{X^\theta}, \ldots, Q_K^\lambda|_{X^\theta}\rangle$$

is a model of φ, φ is valid at $x_1^* \ldots, x_M^*$ for the interpretation of the predicate symbols in w. So $Q_k^\lambda(x_1^* \ldots, x_M^*)$ if and only if $Q_k^*(x_1^* \ldots, x_M^*)$ implies that ϕ is valid at $x_1^* \ldots, x_M^*$

for the interpretation of the predicate symbols in $(R_1^\theta, \ldots, R_N^\theta, Q_1^*|_{X^\theta}, \ldots, Q_K^*|_{X^\theta})$. Thus $(Q_1^*, \ldots, Q_K^*) \in W^\theta$.

The collection W^θ, $\theta \in \Theta$, is thus a nested collection of closed sets in a compact space, so it has nonempty intersection. Conclude that there exists $(Q_1^*, \ldots, Q_K^*) \in \bigcap_{\theta\in\Theta} W^\theta$. We claim that $u^* = (X^*, R_1^*, \ldots, R_N^*, Q_1^*, \ldots, Q_K^*)$ is a model of Σ. Let $\varphi(x_1, \ldots, x_M) \in \Sigma$ and $\{x_1^* \ldots, x_M^*\} \subseteq X^*$. By definition of u^*, there is $\theta \in \Theta$ such that, for $n = 1, \ldots, N$, $R_n^*(x_1^* \ldots, x_M^*)$ if and only if $R_n^\theta(x_1^* \ldots, x_M^*)$, and such that $\{x_1^* \ldots, x_M^*\} \subseteq X^\theta$. Since $(Q_1^*, \ldots, Q_K^*) \in W^\theta$, $\varphi(x_1^*, \ldots, x_M^*)$ is valid under the interpretation $(R_1^\theta, \ldots, R_N^\theta, Q_1^*|_{X^\theta}, \ldots, Q_K^*|_{X^\theta})$. Hence, the fact that $R_n^*(x_1^* \ldots, x_M^*)$ if and only if $R_n^\theta(x_1^* \ldots, x_M^*)$ implies that $\varphi(x_1^*, \ldots, x_M^*)$ is valid under the interpretation $(R_1^*, \ldots, R_N^*, Q_1^*, \ldots, Q_K^*)$. Conclude that $u^* \in T$. As $F(\{u^*\}) = \mathcal{M}^*$, we obtain $\mathcal{M}^* \in F(T)$, establishing the third condition.

Lastly, we now know that $F(T)$ has a universal axiomatization. Obviously, any structure $\mathcal{M} \in F(T)$ satisfies all universal $\mathcal{F}$-implications of Σ. Conversely, we need to show that any sentence in the universal axiomatization of $F(T)$ is a universal $\mathcal{F}$-implication of Σ. So suppose there is a sentence φ which is not. In particular then, there exists a structure $\mathcal{M} \in T$ for which φ is not valid. But as φ involves only predicates from $\mathcal{F}$, it therefore follows that φ is also not valid for $F(\mathcal{M})$, a contradiction (as φ is valid for all members of $F(T)$).

$\square$

2.3 Recursive Axiomatization

In this section, we are relatively informal. Details about computability theory can be found in Rogers (1987).

Falsifiability means that if a theory is incorrect then its incorrectness can be demonstrated. Assume that a scientist believes in some theory T, with universal axiomatization Σ. Suppose that he observes the elements $a_1, \ldots, a_n$ of some structure $\mathcal{M}$, and the relationship between them. If there exist some axiom $\phi(x_1, \ldots, x_n) \in \Sigma$ such that $\phi(a_1, \ldots, a_n)$ is not valid in $\mathcal{M}$ then T has been falsified. On the other hand, if none of the axioms is violated, that does not prove the correctness of the theory, since the axioms may still be violated on other elements of $\mathcal{M}$. This idea is the fundamental tenet of Popper's approach – theories can be falsified, but can never be proved.

We now extend this idea to the process of checking whether observed data (($a_1, \dots, a_n$ and the relations between them) falsify the theory. To do that, our scientist needs an effective procedure to produce a list of all the axioms in Σ, so that he can go over the axioms and check their validity over the data set. Of course, the set of axioms might be infinite (as in the case of the strong axiom of revealed preference, discussed in the next section), which means that if the data violate an axiom, we will eventually find the violations. If none of the axioms are violated, then our search will never end and we will never know for sure that none of the axioms is violated. Again, theories are not proved, only falsified.

Recall that a set Σ of formulas is called *recursively enumerable* (r.e.) if there exists a Turing Machine that enumerates over the elements of Σ in some order. Thus, the output of the machine is an exhaustive list $\phi_1, \phi_2, \dots$ of all the elements of Σ.

The following is a simple corollary of Theorem 5

6 Corollary. *If T has a recursively enumerable and universal axiomatization, then so does $F(T)$*

Proof. Let Σ be a r.e. set of universal formulas that axiomatizes T. It is well known that the set of logical implications of a r.e. set of formulas is itself r.e. Thus, there exists a Turing machine that enumerates all these logical implications $\psi_1, \psi_2, \dots$. We augment this Turing machine by checking before printing each ψ_i whether it is a universal sentence that contains only observable predicates, and print ψ_i only if it satisfies these requirements. The augmented machine enumerates over all universal logical implications of Σ that contain only observable predicates. By Theorem 5, this set axiomatizes $F(T)$. □

7 Remark. A set is *recursive* if both it and its complement are r.e. By a theorem of Craig (1953), a r.e. set Σ of universal axioms is equivalent to a recursive set Σ' of axioms, where each element of Σ' is of the form $\psi \wedge \psi \dots \psi$ for some $\psi \in \Sigma$. Note that the elements of Σ' are not, strictly speaking, universal sentences.

To understand the importance of Corollary 6, consider, for example, the strong axiom of revealed preference. It is, as we emphasize in the next section, an infinite collection of axioms in first order logic. However, there is a simple algorithm that, given a dataset that violates SARP, determines that the theory has been falsified. In Theorem 5 we obtain an axiomatization, but we would like it to be a *test* of the theory, in the same sense as SARP. The property of recursive enumerability ensures that an axiomatization is such a test.

Another interpretation of Corollary 6 relies on an equivalent definition of recursive enumerability: A set Σ is r.e. if there exists a Turing Machine that *recognizes* Σ, i.e. halts on an input ϕ if and only if $\phi \in \Sigma$. The Turing machine that recognizes Σ can be thought of as a method of proof: to prove that some formula ϕ is in Σ, one needs to run the machine until it halts. (On the other hand, if Σ is r.e. and $\phi \notin \Sigma$, then this assertion may not be provable). Thus, Σ is r.e. if membership in Σ, if true, is provable.

Assume that a theory T has a universal r.e. axiomatization Σ, then if a structure $\mathcal{M}$ is not in T then this fact is demonstrable in the following way: First, one has to come up with an element $\phi(x_1, \ldots, x_n)$ of Σ and elements $a_1, \ldots, a_n$ of M such that $\phi(a_1, \ldots, a_n)$ is not valid in $\mathcal{M}$. Second, one has to use the Turing machine that recognizes Σ to prove that indeed $\phi \in \Sigma$. Thus, the existence of a universal r.e. axiomatization means that if a structure is not in the theory then this fact can be demonstrated. This is what falsifiability of T means.

2.4 Example: Individual rational choice

As an illustration of Theorem 5, consider the revealed preference formulation of the theory of individual rational choice.

Consider the language $\mathcal{F} = \langle R, P \rangle$ with two binary predicates; $R(x, y)$ is intended to mean that x is revealed preferred to y, and $P(x, y)$ that x is revealed strictly preferred to y.

We are interested in the theory of all structures (X, R^X, P^X) for which there exists a complete and transitive binary relation $\succeq$ satisfying the axioms

1. $\forall x \forall y (R(x, y) \to \succeq (x, y))$
2. $\forall x \forall y (P(x, y) \to \succ (x, y))$

The language $\mathcal{F}$ expresses only observables, but the description we have of the theory does not constitute a universal axiomatization. We want to know when we can formulate the theory using only universal axioms that are statements about observables.

We can extend the language to a language $\mathcal{L}$, and formulate our theory using only universal axioms. Let $\mathcal{L} = \langle R, P, \succeq, \succ \rangle$. Consider the set of $\mathcal{L}$-sentences:

1. $\forall x \forall y (\succeq (x, y) \vee \succeq (y, x))$
2. $\forall x \forall y (\succ (x, y) \leftrightarrow (\succeq (x, y) \wedge \neg \succeq (y, x)))$
3. $\forall x \forall y \forall z (\succeq (x, y) \wedge \succeq (y, z)) \rightarrow \succeq (x, z)$
4. $\forall x \forall y (R(x, y) \rightarrow \succeq (x, y))$
5. $\forall x \forall y (P(x, y) \rightarrow \succ (x, y))$

Let T be the $\mathcal{L}$-theory axiomatized by this set of sentences.

Now, T is a description of the theory of rational choice, but it assumes that we can access, or observe, the relation $\succeq$. The theory we are really interested in is $F(T)$: the projection of T onto the language $\mathcal{F}$.

Since the axioms (1)-(5) are a universal axiomatization of T, Theorem 5 implies that there is a universal axiomatization of $F(T)$. In addition, one such axiomatization is given by the implications of (1)-(5) that only involve the predicates R and P.

To be concrete, this axiomatization is described by a variant of the strong axiom of revealed preference. In fact, in first order logic, the strong axiom of revealed preference is a collection of axioms.

The Strong Axiom of Revealed Preference: For every k,

$$\forall x_1 ... \forall x_k \neg \bigwedge_{i=1}^{k} \left(x_i \, Q_i \, x_{(i+1) \mod k} \right)$$

where for all i, $Q_i \in \{R, P\}$, and for at least one $i \in \{1, ..., k\}$, $Q_i = P$.

8 Proposition. *If T is axiomatized by (1)-(5), then $F(T)$ is axiomatized by the strong axiom of revealed preference.*

Proof. We offer a sketch, as this type of argument is well-understood. Clearly the strong axiom is valid for $F(T)$. Now suppose that (X, R^X, P^X) is a model of the theory described by the strong axiom. We want to show it is an element of $F(T)$. Let Q denote the transitive closure of $R^X \cup P^X$. Note that if $P^X(x, y)$, then $\neg Q(y, x)$ (this follows by the strong axiom of revealed preference). Consequently, denoting the strict part of Q by P_Q, we obtain $\forall x \forall y R^X(x,y) \rightarrow Q(x,y)$ and $\forall x \forall y P^X(x,y) \rightarrow P_Q(x,y)$. Now, by a generalization of the Szpilrajn Theorem (see, for example, Suzumura (1976), Theorem 3), Q has an extension to a weak order $\succeq$ with strict part $\succ$, so that $\forall x \forall y Q(x,y) \rightarrow \succeq (x,y)$ and $\forall x \forall y P_Q(x,y) \rightarrow \succ (x,y)$. Consequently, $\forall x \forall y R^X(x,y) \rightarrow \succeq (x,y)$ and $\forall x \forall y P^X(x,y) \rightarrow \succ (x,y)$, where $\succeq$ is a weak order and $\succ$ is its strict part. This verifies that $(X, R^X, P^X) \in F(T)$, as $(X, R^X, P^X, \succeq, \succ) \in T$. $\square$

3 Rationalizing group choice behavior

In this section, we look at Nash equilibrium behavior. We assume that we observe a collection of game forms, and a choice made from each game form. We ask whether or not there could exist *strict* preferences for a collection of agents over those game forms which generate the observed choices as Nash equilibrium behavior. We show, using Theorem 5, that this theory has a universal axiomatization.

We first have to set up our framework. Instead of focusing on Nash equilibrium specifically, we work with a general collection of theories of group choice. Nash equilibrium, strong Nash equilibrium, and Pareto optimal choice are special cases. We fix a finite set of agents $N = \{1, ..., n\}$ and a collection $\Gamma \subseteq 2^N \backslash \{\varnothing\}$. The elements of Γ are the sets of agents that can deviate from a profile of strategies.

A *game form* is a tuple $(S_1, ..., S_n)$, where we think of S_i as the set of strategies available to agent i. For each profile of preferences $(\succ_1, \ldots, \succ_n)$, a game form $(S_1, ..., S_n)$ defines a normal-form game

$$(S_1, ..., S_n, \succ_1, \ldots, \succ_n).$$

Here preferences $\succ_j$ are defined over $\prod_{i \in N} S_i$, the set of all strategy profiles.

We will define a Γ-Nash equilibrium of a game $(S_1, ..., S_n, \succ_1, ..., \succ_n)$ to be $s \in \prod_{i \in N} S_i$ for which for all $\gamma \in \Gamma$ and all $s' \in \prod_{i \in N} S_i$, if there exists $j \in \gamma$ for which $(s'_\gamma, s_{-\gamma}) \succ_j s$, then there exists $k \in \gamma$ for which $s \succ_k (s'_\gamma, s_{-\gamma})$.

So the following are special cases:

- Nash equilibrium results when $\Gamma = \{\{i\} : i \in N\}$
- Pareto optimality results when $\Gamma = \{N\}$

- Strong Nash equilibrium results when $\Gamma = 2^N \backslash \{\varnothing\}$

Other kinds of theories are permissible. For example, by setting $\Gamma = \{G : |G| > |N|/2\}$, we get a kind of majority rule core.

We imagine that we observe a collection of game forms, and *some* strategy profiles which are chosen from each. We do not necessarily observe the entire collection of strategy profiles which could potentially be chosen.

We ask when the strategy profiles are rationalizable by a list of preference relations; obviously, if we make no restriction on preferences, then *every* strategy profile is rationalizable by complete indifference. To this end, we require that preferences be *strict* over strategy profiles. This is a significant assumption.

Let us define the *language of group choice* $\mathcal{F}$ to include the following predicates:

- For each $i \in N$, one unary predicate S_i, where $S_i(y)$ is intended to mean that y is a set of strategies for i
- For each $i \in N$, one unary predicate s_i, where $s_i(x)$ means that x is a strategy for i
- The typical set theoretic binary predicate $\in$, meant to signify participation in a set
- A $2n$-ary predicate R, where $R(y_1, ..., y_n, x_1, ..., x_n)$ means that $(x_1, ..., x_n)$ is observed as being chosen from game form $(y_1, ..., y_n)$

The *theory of group choice* T_G is the class of all structures for the preceding language for which there exists for each agent a *global strategy space* $\mathcal{S}_i$ for which the following objects are elements of the universe

- Each nonempty $S_i \subseteq \mathcal{S}_i$
- Each $s_i \in \mathcal{S}_i$

and each of the predicates S_i, s_i, $\in$ are interpreted properly. For each game form $\prod_i S_i^*$ and strategy profile $(s_1^*, ..., s_n^*)$, $R(S_1^*, ..., S_n^*, s_1^*, ..., s_n^*)$ implies that $S_i(S_i^*)$, $s_i(s_i^*)$, and lastly that $s_i^* \in S_i^*$.

The *theory of* Γ*-rationalizable choice* $T_\Gamma \subseteq T_G$ is the theory of group choice for which for each $i \in N$, there exists a linear order $\succ_i$ over $\prod_{i\in N} \mathcal{S}_i$ for which for all game forms $\prod_i S_i^*$, $R(S_1^*, ..., S_n^*, s_1^*, ..., s_n^*)$ implies that $(s_1^*, ..., s_n^*)$ is a Γ-Nash equilibrium of the normal-form game $(S_1^*, ..., S_n^*, \succ_1, \ldots, \succ_n)$.

9 Theorem. *The theory of* Γ*-rationalizable choice is universally and recursively enumerably axiomatizable with respect to the theory of group choice.*

Note that Theorem 9 deals with the universal axiomatization of Γ-rationalizable choice with respect to the theory of group choice. We do not here want to focus on axiomatizing group choice; we want to focus only on the additional empirical content of Γ-Nash equilibrium.

Proof. Consider the language $\mathcal{L}$ which includes all predicates in $\mathcal{F}$, but also includes, for each agent i, a $2n$-ary predicate $\succ_i$.

We use the abbreviation

$$\succeq_k (x_1, ..., x_n, z_1, ..., z_n) = \succ_k (x_1, ..., x_n, z_1, ..., z_n) \vee \left(\bigwedge_{i=1}^{n} x_i = z_i \right).$$

For each $\gamma \in \Gamma$ and $k \in \gamma$, we use the following shorthand:

If $|\gamma| > 1$,

$$\begin{aligned} &\varphi_{\gamma,k}(x_1, ..., x_n, z_1, ..., z_n) = \\ \succ_k ((z_\gamma, x_{-\gamma}), x) \rightarrow &\bigvee_{i \in \gamma \setminus \{k\}} \succ_i (x, (z_\gamma, x_{-\gamma})) \end{aligned}$$

otherwise, if $|\gamma| = 1$,

$$\varphi_{\gamma,k}(x_1, ..., x_n, z_1, ..., z_n) = \succeq_k (x_1, ..., x_n, z_1, ..., z_n).$$

Consider the theory T axiomatized by the following sentences.

For each $\gamma \in \Gamma$ and $k \in \gamma$,

$$\forall x_1 ... \forall x_n \forall y_1 ... \forall y_n \forall z_1 ... \forall z_n$$
$$\bigwedge_{i\in\gamma} \in (z_i, y_i) \wedge \bigwedge_{i\in N} \in (x_i, y_i) \wedge R(y_1, ..., y_n, x_1, ..., x_n)$$
$$\rightarrow \varphi_{\gamma,k}(x_1, ..., x_n, z_1, ..., z_n)$$

Completeness: For each $k \in N$,

$$\forall s_1 ... \forall s_n \forall s'_1 ... \forall s'_n$$
$$\bigwedge_{i=1}^{n} s_i(s_i) \wedge \bigwedge_{i=1}^{n} s_i(s'_i) \rightarrow (\succ_k (s_1, ..., s_n, s'_1, ..., s'_n) \vee \succ_k (s'_1, ..., s'_n, s_1, ..., s_n))$$

Transitivity:

$$\forall s_1 ... \forall s_n \forall s'_1 ... \forall s'_n \forall s''_1 ... \forall s''_n$$
$$\bigwedge_{i=1}^{n} s_i(s_i) \wedge \bigwedge_{i=1}^{n} s_i(s'_i) \wedge \bigwedge_{i=1}^{n} s_i(s''_i)$$
$$\rightarrow (\succ_k (s_1, ..., s_n, s'_1, ..., s'_n) \wedge \succ_k (s'_1, ..., s'_n, s''_1, ..., s''_n) \rightarrow \succ_k (s_1, ..., s_n, s''_1, ..., s''_n))$$

Asymmetry:

$$\forall s_1 \forall s_2 ... \forall s_n \forall s'_1 ... \forall s'_n$$
$$\left(\bigwedge_{i=1}^{n} s_i(s_i) \wedge \bigwedge_{i=1}^{n} s_i(s'_i)\right) \vee \left(\bigvee_{i=1}^{n} s_i \neq s'_i\right)$$
$$\neg (\succ_k (s_1, ..., s_n, s'_1, ..., s'_n) \wedge \succ_k (s'_1, ..., s'_n, s_1, ..., s_n))$$

As T has a universal axiomatization, so does $F(T)$. Since the axiomatization of T is finite, $F(T)$ has a recursively enumerable universal axiom by Corollary 6. And $T_\Gamma = F(T) \cap T_G$. So T_Γ has a r.e. universal axiomatization within T_G.

□

Because we ask for group choice functions to be rationalizable by strict preferences, the exercise here is similar in spirit to Afriat (1967), who assumes partial observations of demand functions and asks demand functions to be rationalizable by locally non-satiated preference.

On the other hand, we could construct an exercise similar to Richter (1966), assuming that not all game forms are observable, but that given a game form is observed, *all* possible choices from that game form are observed. Studying such an object would involve constructing a richer language (that allows us to speak of game forms), but we also conjecture that the theory of Γ-rationalizable group choice functions is *not* universally axiomatizable in this context.[3]

4 Relation to the existing previous literature

The type of issues we discuss here have previously been studied by philosophers of science. Without going into full detail, Ramsey (1931) was one of the first to discuss the elimination of "theoretical" terms from scientific theories. Various authors give different interpretation to the notion of "Ramsey elimination." The work of Sneed (1971) includes notions very similar to ours; in particular, he defines a finitely axiomatized $\mathcal{L}$- theory T to be Ramsey eliminable if $F(T)$ is a finitely axiomatized $\mathcal{F}$-theory. In particular, he includes an example (attributed to Dana Scott) of a theory T which is first order axiomatizable, but for which $F(T)$ is not first order axiomatizable. We include here an adaptation of this example.

10 Example. Let $\mathcal{F} = \langle R \rangle$, where R is a unary predicate, and let $\mathcal{L} = \langle R, Q \rangle$, where Q is a binary predicate. Consider the $\mathcal{L}$-theory T axiomatized by the following sentences:

1. $\forall x \forall y Q(x, y) \rightarrow R(x) \wedge \neg R(y)$
2. $\forall x R(x) \rightarrow (\exists y Q(x, y) \wedge (\forall z Q(x, z) \rightarrow y = z))$
3. $\forall x \neg R(x) \rightarrow (\exists y Q(y, x) \wedge (\forall z Q(z, x) \rightarrow y = z))$

[3] This departs from the single agent case, where single agent choice is known to be universally rationalizable.

Thus, T is the theory of all structures (X, R^X, Q^X) for which there is a one-to-one correspondence Q^X between the elements of R^X and its complement.

11 Proposition. *$F(T)$ is not first order $\mathcal{F}$-axiomatizable.*

Proof. Suppose by means of contradiction that there is a first order axiomatization of $F(T)$. Consider a structure (X, R^X) where $|R^X|$ is infinite, $|X \backslash R^X|$ is infinite, and the cardinalities of R^X and $X \backslash R^X$ are distinct. Note that $(X, R^X) \notin F(T)$.

By the Löwenheim-Skolem theorem, (see for example Marker (2002), Theorem 2.3.7), there exists a countable structure $(X', R^{X'})$ which satisfies exactly the same first order sentences as (X, R^X). But note in particular that for each $n > 0$, the sentence

$$\exists x_1 ... \exists x_n \bigwedge_{i=1}^{n} R(x_i) \wedge \bigwedge_{i \neq j} (\neg(x_i = x_j))$$

is valid for (X, R^X); in particular, then, it is valid for $(X', R^{X'})$; consequently, $|R^{X'}|$ is infinite. Similarly, since

$$\exists x_1 ... \exists x_n \bigwedge_{i=1}^{n} \neg R(x_i) \wedge \bigwedge_{i \neq j} (\neg(x_i = x_j))$$

is valid for (X, R^X), it is also valid for $(X', R^{X'})$ and hence $|X' \backslash R^{X'}|$ is infinite. Since X' is countable, there is therefore a bijection between $R^{X'}$ and $X \backslash R^{X'}$, so that $(X', R^{X'}) \in F(T)$. But then (X, R^X) satisfies the sentences axiomatizing $F(T)$ (as it satisfies the same sentences as $(X', R^{X'})$ and $(X', R^{X'}) \in F(T)$). So $(X, R^X) \in F(T)$, a contradiction. $\square$

The preceding example is important in that it illustrates that the problem of axiomatizability of $F(T)$ is non-trivial. To this end, Van Benthem (1978) (Theorem 4.2) uncovers necessary and sufficient conditions for $F(T)$ to be first order axiomatizable, given that T is first order axiomatizable. His conditions are essentially an adaptation of a well-known theorem in model theory axiomatizing those theories which are first order axiomatizable (see, for example Chang and Keisler (1990) Theorem 4.1.12). The condition is a semantic condition requiring one to verify, for any model of $F(T)$, whether any structure which satisfies exactly the same $\mathcal{F}$-sentences is also a model of $F(T)$. In practice, this is nearly impossible to verify. Our condition; on the other hand, is a syntactic condition which is trivial to verify.

Economists have also discussed these issues. It is well-known to economists that the theory of rational choice is falsifiable. Simon (1985) discusses the issue we discuss here. Simon argues that the theory is falsifiable, even though the RP formulation of the theory is existentially quantified over unobservables. As Simon (1985) states, "although existential quantification of an observable is fatal to the falsifiability of a theory, the same is not true when the existentially quantified term is a theoretical one."

While this may seem obvious to many, it has led to a large degree of confusion among economists. For example Boland (1981) argued that the theory of rational choice is not falsifiable precisely because of its existential formulation over unobservables. In his words:

> Given the premise–"All consumers maximize something"–the critic can claim he has found a consumer who is not maximizing anything. The person who assumed the premise is true can respond: "You claim you have found a consumer who is not a maximizer but how do you know there is not something which he is maximizing?" In other words, the verification of the counterexample requires the refutation of a strictly existential statement; and as stated above, we all agree that one cannot refute existential statements.

Mongin (1986) beautifully counters this argument. In this context, he has already observed that all $\mathcal{F}$-implications of T are $\mathcal{F}$-implications of $F(T)$. It follows from this that if T is universal, then $F(T)$ has universal implications, and is hence falsifiable. As noted above, we have gone further than this in showing that in fact, $F(T)$ is first-order axiomatizable (indeed, universally axiomatizable). Hence *all* of its implications are falsifiable.

Our work is related to the approach in Brown and Matzkin (1996), and the general approach to testable implications discussed in Brown and Kubler (2008). In these papers, as in ours, there is an operation of projection to eliminate certain existential quantifications. The idea in Brown-Matzkin and in Brown-Kubler is to exploit the property of quantifier elimination in certain mathematical theories. Our work, on the other hand, uses results from model theory on when a universal axiomatization is possible. Our projection argument follows from the verification that the universal axiomatization can be projected. We do not exploit the property of quantifier elimination of the mathematical theory underlying the economic model (indeed our results may apply when quantifier elimination does not hold).

In our previous paper, Chambers, Echenique, and Shmaya (2010), we dealt with theories which could be axiomatized by what we called *UNCAF* formulas, for universal negation of conjunction of atomic formulas. Under certain conditions, being UNCAF axiomatizable is equivalent to being universally axiomatizable. And in fact, we show there that a result akin to Theorem 5 is true for UNCAF sentences. This proof relies on drastically different techniques, as there are deep differences between theories axiomatizable by UNCAF sentences and those which are only universally axiomatizable. In particular, theories which are UNCAF axiomatizable are closed under weak substructures (and not just substructures), a property that plays a critical role in our previous paper.

The reason Theorem 5 is useful is because, in general, the hypothesized theory T is usually not axiomatizable by UNCAF sentences. For example, the axioms of completeness and transitivity for binary relations have no UNCAF axiomatization. In general, then, it is impossible to obtain any results about the falsifiability of a revealed preference theory without having some result about universality.

Appendix A Basic definitions from Model Theory

The following definitions are taken, for the most part, quite literally from (Marker, 2002), pp. 8-12. We refer readers to this excellent text for more details; but present the basics here to keep the analysis self-contained. The $\overline{x}$ notation is here used to denote a list, or vector, or elements $(x_1, ..., x_m)$.

12 Definition. $\mathcal{M}$ is a *substructure* of $\mathcal{N}$ if $M \subseteq N$ and the inclusion map $\iota : M \to N$ defined by $\iota(m) = m$ for all $m \in M$ is an $\mathcal{L}$-embedding.

The following definition gives us the basic building blocks of our syntax. Note that we include a countable list of "variables" to be used in this definition; these are not part of the language *per se*, but rather part of a "meta language" in that they are present in all languages.

13 Definition. The set of $\mathcal{L}$*-terms* is the smallest set $\mathcal{TE}$ such that

1. $c \in \mathcal{TE}$ for each constant symbol $c \in \mathcal{C}$
2. each variable symbol $v_i \in \mathcal{TE}$ for $i = 1, 2, ...,$
3. if $t_1, ..., t_{n_f} \in \mathcal{TE}$ and $f \in \mathcal{F}$, then $f(t_1, ..., t_{n_f}) \in \mathcal{TE}$.

The following definitions mark our departure from Marker. Specifically, we want to allow atomic formulas to include expressions involving the $\neq$ sign–and we want to include this symbol as part of our meta-language, in the sense that it is present in every language.

14 Definition. Say that ϕ is an *atomic* $\mathcal{L}$*-formula* if ϕ is one of the following

1. $t_1 = t_2$, where t_1 and t_2 are terms
2. $t_1 \neq t_2$, where t_1 and t_2 are terms
3. $R(t_1, ..., t_{n_R})$, where $R \in \mathcal{R}$ and $t_1, ..., t_{n_R}$ are terms

15 Definition. The set of $\mathcal{L}$*-formulas* is the smallest set $\mathcal{W}$ containing the atomic formulas such that

1. if ϕ is in $\mathcal{W}$, then $\neg\phi$ is in $\mathcal{W}$
2. if ϕ and ψ, then $(\phi \wedge \psi)$ and $(\phi \vee \psi)$ are in $\mathcal{W}$
3. if ϕ is in $\mathcal{W}$, then $\exists v_i \phi$ and $\forall v_i \phi$ are in $\mathcal{W}$.

16 Definition. A variable v *occurs freely* in a formula ϕ if it is not inside a $\exists v$ or $\forall v$ quantifier. It is *bound* in ϕ if it does not occur freely in ϕ.

17 Definition. A *sentence* is a formula ϕ with no free variables.

We are now prepared to define a concept of "truth" relating syntax and semantics. We want to define what it means for a sentence to be true in a given structure. The notion we define here is slightly different than Marker, as it again relies on the correct interpretation of the $\neq$ symbol, which is not a primitive there (nor in any other standard text).

18 Definition. Let ϕ be a formula with free variables from $\overline{v} = (v_{i_1}, ..., v_{i_m})$, and let $\overline{a} = (a_{i_1}, ..., a_{i_m}) \in M^m$. We inductively define $M \models \phi(\overline{a})$ as follows. The notation $M \not\models \psi(\overline{a})$ means that $M \models \phi(\overline{a})$ is not true.

1. If ϕ is $t_1 = t_2$, then $\mathcal{M} \models \phi(\overline{a})$ if $t_1^{\mathcal{M}}(\overline{a}) = t_2^{\mathcal{M}}(\overline{a})$
2. If ϕ is $t_1 \neq t_2$, then $\mathcal{M} \models \phi(\overline{a})$ if $t_1^{\mathcal{M}}(\overline{a}) \neq t_2^{\mathcal{M}}(\overline{a})$
3. If ϕ is $R(t_1, ..., t_{n_R})$, then $\mathcal{M} \models \phi(\overline{a})$ if $(t_1^{\mathcal{M}}(\overline{a}), ..., t_{n_R}^{\mathcal{M}}(\overline{a})) \in R^{\mathcal{M}}$
4. If ϕ is $\neg\psi$, then $\mathcal{M} \models \phi(\overline{a})$ if $\mathcal{M} \not\models \psi(\overline{a})$
5. If ϕ is $(\psi \wedge \theta)$, then $\mathcal{M} \models \phi(\overline{a})$ if $\mathcal{M} \models \psi(\overline{a})$ and $\mathcal{M} \models \theta(\overline{a})$
6. If ϕ is $(\psi \vee \theta)$, then $\mathcal{M} \models \phi(\overline{a})$ if $\mathcal{M} \models \psi(\overline{a})$ or $\mathcal{M} \models \theta(\overline{a})$
7. If ϕ is $\exists v_j \psi(\overline{v}, v_j)$, then $\mathcal{M} \models \phi(\overline{a})$ if there is $b \in M$ such that $\mathcal{M} \models \psi(\overline{a}, b)$
8. If ϕ is $\forall v_j \psi(\overline{v}, v_j)$, then $\mathcal{M} \models \phi(\overline{a})$ if for all $b \in M$, $\mathcal{M} \models \psi(\overline{a}, b)$.

19 Definition. $\mathcal{M}$ *satisfies* $\phi(\overline{a})$ or $\phi(\overline{a})$ *is true* in $\mathcal{M}$ if $\mathcal{M} \models \phi(\overline{a})$.

Lastly, for our purposes, it is useful to have a notion of a *universal* sentence.

20 Definition. A *universal sentence* or *universal formula* is a sentence of the form $\forall \overline{v} \phi(\overline{v})$, where ϕ is quantifier free.

References

Afriat, S. N. (1967): "The Construction of Utility Functions from Expenditure Data," *International Economic Review*, 8(1), 67–77.

Aumann, R. (1960): "Acceptable points in games of perfect information," *Pacific Journal of Mathematics*, 10(2), 381–417.

Boland, L. (1981): "On the futility of criticizing the neoclassical maximization hypothesis," *The American Economic Review*, 71(5), 1031–1036.

Brown, D., and F. Kubler (2008): "Refutable Theories of Value," in *Computational Aspects of General Equilibrium Theory: Refutable Theories of Value*. Springer Publishing Company, Incorporated.

Brown, D. J., and R. L. Matzkin (1996): "Testable Restrictions on the Equilibrium Manifold," *Econometrica*, 64(6), 1249–1262.

Chambers, C. P., F. Echenique, and E. Shmaya (2010): "The axiomatic structure of empirical content," working paper.

Chang, C., and H. Keisler (1990): *Model theory*. North Holland.

Craig, W. (1953): "On Axiomatizability Within a System," *The Journal of Symbolic Logic*, 18(1), 30–32.

Galambos, A. (2009): "The complexity of Nash rationalizability," working paper.

Houthakker, H. (1950): "Revealed preference and the utility function," *Economica*, pp. 159–174.

Ledyard, J. O. (1986): "The scope of the hypothesis of Bayesian equilibrium," *Journal of Economic Theory*, 39(1), 59 – 82.

Lee, S. (2009): "The testable implications of zero-sum games," Mimeo, California Institute of Technology.

Marker, D. (2002): *Model theory: an introduction*. Springer Verlag.

Mongin, P. (1986): "Are All-and-Some Statements Falsifiable After All?," *Economics and Philosophy*, 2, 185–195.

Plott, C. R. (1974): "On Game Solutions and Revealed Preference Theory," Discussion Paper 35, California Institute of Technology.

Popper, K. R. (1959): *The Logic of Scientific Discovery*. Hutchinson.

Ramsey, F. (1931): "Theories," *The Foundations of Mathematics*, pp. 212–236.

Richter, M. K. (1966): "Revealed Preference Theory," *Econometrica*, 34(3), 635–645.

Rogers, H. (1987): *Theory of Recursive Functions and Effective Computability*. MIT Press, Cambridge, MA.

Samuelson, P. (1938): "A note on the pure theory of consumer's behaviour," *Economica*, pp. 61–71.

Simon, H. (1985): "Quantification of theoretical terms and the falsifiability of theories," *The British Journal for the Philosophy of Science*, 36(3), 291.

Sneed, J. (1971): *The Logical Structure of Mathematical Physics*. D. Reidel Publishing Company.

Sprumont, Y. (2000): "On the testable implications of collective choice theories," *Journal of Economic Theory*, 93(2), 205–232.

Suzumura, K. (1976): "Remarks on the Theory of Collective Choice," *Economica*, 43(172), 381–390.

Tarski, A. (1954): "Contributions to the theory of models I," *Indagationes Mathematicae*, 16, 572–581.

Van Benthem, J. (1978): "Ramsey eliminability," *Studia Logica*, 37(4), 321–336.

Wilson, R. B. (1970): "The finer structure of revealed preference," *Journal of Economic Theory*, 2(4), 348 – 353.

Xu, Y., and L. Zhou (2007): "Rationalizability of choice functions by game trees," *Journal of Economic Theory*, 134(1), 548–556.