DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

THE AXIOMATIC STRUCTURE OF EMPIRICAL CONTENT

Christopher P. Chambers
California Institute of Technology

Federico Echenique
California Institute of Technology

Eran Shmaya
Northwestern University

SOCIAL SCIENCE WORKING PAPER 1317

January 2010

The Axiomatic Structure of Empirical Content

Christopher P. Chambers Federico Echenique Eran Shmaya

Abstract

In this paper, we provide a formal framework for studying the empirical content of a given theory. We define the **falsifiable closure** of a theory to be the least weakening of the theory that makes only falsifiable claims. The falsifiable closure is our notion of empirical content. We prove that the empirical content of a theory can be exactly captured by a certain kind of axiomatization, one that uses axioms which are universal negations of conjunctions of atomic formulas. The falsifiable closure operator has the structure of a topological closure, which has implications, for example, for the behavior of joint vis a vis single hypotheses.

The ideas here are useful for understanding theories whose empirical content is well-understood (for example, we apply our framework to revealed preference theory, and Afriat's theorem), but they can also be applied to theories with no known axiomatization. We present an application to the theory of multiple selves, with a fixed finite set of selves and where selves are aggregated according to a neutral rule satisfying independence of irrelevant alternatives. We show that multiple selves theories are fully falsifiable, in the sense that they are equivalent to their empirical content.

JEL classification numbers: A10,D00

Key words: Falsifiability; Empirical Content ; Revealed preference; Afriat's Theorem; First-order logic; Model Theory

The Axiomatic Structure of Empirical Content*

Christopher P. Chambers   Federico Echenique   Eran Shmaya

1 Introduction

Falsifiability has been a hallmark of the scientific method at least since Popper (1959). The predictive power of a theory is only as good as the falsifiable claims that it makes. Any two theories making the same falsifiable claims are observationally equivalent. This paper is an axiomatic study of the empirical content of a theory.

Most economic theories make some falsifiable claims, but not all its claims may be falsifiable. An example that would be familiar to most economists is the theory of utility maximization. If we imagine that we can observe data on choices among pairs, then we can test for the transitivity of preference: transitivity is a testable implication of the theory of utility maximization. But we cannot test for the existence of a rationalizing utility. In fact, it turns out that the theory of utility maximization has the same empirical content as the theory of preference maximization.[1] In the terminology of our paper, the theory of preference maximization is the *falsifiable closure* of the theory of utility maximization. The falsifiable closure is our notion of empirical content.

Some theories make only falsifiable claims. Such theories are called *falsifiably complete*; one example is the theory of preference maximization (although this depends on the possible data taken as primitive, as we explain below). A theory is falsifiably complete if all claims it makes can be refuted empirically. One can weaken a theory by, in a sense, discarding some of its claims, until it becomes falsifiably complete. The falsifiable

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

[1]Our discussion in the introduction is necessarily very loose. Example 3 presents these theories more rigorously.

closure of a theory is the *least* such weakening. A theory is falsifiably complete if it coincides with its falsifiable closure; eliminating any of its claims results in an observationally distinct theory.

Our main result is a syntactic characterization of falsifiable closure, and of falsifiably complete theories. This allows us to understand the syntax (the formal structure) of axioms that ensure falsifiable completeness, and that characterize the empirical content of a theory. We apply our results to theories from behavioral economics and social choice; theories that were not previously known to be fully testable.

The data that can be observed is a primitive of our model. For example, consider again the theory of preference maximization. It matters whether we believe that we can observe weak preference only (an alternative x is weakly preferred to y), or if we can also observe strict preference. It also matters whether we can observe absence of preference: do we allow observations of the type "x is not preferred to y?"

We restrict the possible data by assuming a language. For example, if we want to assume that we can observe preference and absence of preference, then we can assume that we have two symbols R and $\tilde{R}$; we intend to use R to denote a binary relation expressing weak preference, and $\tilde{R}$ to denote absence of weak preference.[2] We may hypothesize the seemingly tautological statement that between any pair, there is either preference or absence of preference, and never both. This illustrates the role language plays in our framework. The key here is that our language needs to be rich enough to allow us to discern between absence of an observation of preference, and the observation of absence of preference.

Given a language, we can write axioms expressed in the symbols of the language. These are the first-order sentences that can be expressed in the language, a notion from mathematical logic. An axiomatization of a theory is a collection of sentences which hold at, and only at, each of the particular instances of the theory. We prove that axioms that have a certain form (universal negations of conjunctions of atomic formulas, or UNCAF) characterize the empirical content of a theory. So the falsifiable closure of a theory is axiomatized by UNCAF sentences.

For example, if we use the language with the symbols R and $\tilde{R}$, an UNCAF axiom is

[2]This is similar to being able to observe both preference and strict preference; however, for formal exposition, these turn out to be the most tractable primitives, for reasons that will become apparent.

the following sentence:

$$\forall x \forall y \forall z \neg[(x\, R\, y) \wedge (y\, R\, z) \wedge (z\, \tilde{R}\, x)];$$

an axiom expressing transitivity.

We show that the falsifiable closure is a topological closure operator on theories. This is important, as it implies that the intersection of an arbitrary number of falsifiably complete theories is also falsifiably complete. In other words, no new implications can be derived from the intersection of the theories; other than those which follow as logical consequences of the implications already present. On the other hand, the empirical content of joint hypotheses may be strictly more (i.e. imply stronger restrictions) than the intersection of the content of each individual hypothesis, but only when these individual theories are not falsifiably complete.

Our paper formalizes existing ideas and notions. The formalization raises new and subtle issues, and our results are readily applicable to economic theories. First, we demonstrate this by characterizing the empirical content of a large body of theories which as of yet have resisted axiomatization: the theory of multiple selves. We model the multiple selves hypothesis by assuming a fixed and finite set of selves, each of whom has a strict preference. The preferences of these selves are aggregated according to a preference aggregation rule satisfying two very simple hypotheses (neutrality and IIA). We in fact demonstrate that all such models are equivalent to their empirical content: they are falsifiably complete, and thus possess an UNCAF axiomatization. This is by no means a trivial or well-known result, as a special case of it is closely related to the unsolved dimension problem of order theory: our result shows that the class of orders having dimension less than n for any fixed n has a very special type of axiomatization; we believe this result to be previously unknown (see Trotter (1992)).

Secondly, while the results we discuss seem intuitive and many might feel they are already familiar with the ideas, there are many subtleties involved. In fact, no results can be true without hypotheses; and in fact, uncovering hypotheses which render the result true is, in this case, a nontrivial task. For example; Popper (1959) essentially regarded falsifiable theories and universal theories as equivalent.[3] We have shown that universality of a theory is not in general strong enough to imply falsifiable completeness. We hope that our formal presentation will aid in understanding the appropriate relationship between

[3] A universal theory is one which has a universal axiomatization; relying only on universal quantification

observable data, hypotheses, and the axiomatic approach, at least for positive economics.

We do not wish to suggest that theories which are not falsifiably complete are not economically interesting. Many normative studies recommend situations which cannot be empirically falsified. An example of an interesting normative theory which is not falsifiable is the theory of *egalitarian-equivalence* of Pazner and Schmeidler (1978).[4] Egalitarian equivalence is a theory of private goods consumption which requires that there exists some reference bundle for which everyone is indifferent between her own consumption and the reference bundle. Such a requirement is clearly meant to be prescriptive, rather than descriptive.

1.1 The basic formal idea

To study the structure of axioms, we need to have a way of talking about axioms as formal mathematical concepts. The mathematical field of model theory provides us with tools for such an analysis. Our paper uses definitions and basic ideas from model theory. We model the data we can observe by a first order language, involving relation and function symbols. The relation and function symbols should be chosen to correspond to things we think of as primitive observables. For example if we believe we can observe a preference between a pair of alternatives, we need to include a relation for that preference; and if we believe we can observe the absence of preference, we need to include a *separate* relation corresponding to absence of preference.

When we speak of a finite data set, we mean a finite set equipped with relations (corresponding to the relations) and partial functions (a function whose domain is a subset of the finite set–corresponding to the function symbols of our language). The data set involves all elements we have observed to stand in a given relation, as well as the image of certain elements of our domain under certain functions. In the preference example, we have no function symbols, and a data set is simply a set with a pair of binary relations (relating to weak preference and absence of weak preference). The binary relations represent all observed comparisons. The example illustrates a critical assumption underlying our work: we do not want to equate non-observation of a relation with the observation of a negation of the relation. In the example, non observation of x being weakly preferred to y does not mean we have observed x is not weakly preferred to y.

[4]This theory though, is easily *verified*; see Appendix A.

We define a theory as a class of structures for the language, which is closed under isomorphism. The idea is that a theory postulates a class of structures which we believe may represent the real world. We do not know which, if any, of these structures actually represents the real world; this will be something we have to test. That theories must be closed under isomorphism is postulated for technical reasons, but it is fairly intuitive that a class of structures which is not closed under isomorphism would be strange indeed. The reason is that, given our language, we can only describe how certain things interact, but we cannot refer to specific structures themselves. Any two structures whose objects interact in the same way will be indistinguishable. This notion of theory is therefore as general as could possibly be while still being a meaningful concept. Theories will be denoted T and T'.

Because as economists we often want to assume that a certain theory of behavior is implicit, we provide relative definitions of falsifiability. Our main results are a characterization of falsifiably complete theories and characterizations of falsifiable closure (as well as properties of the falsifiable closure operator). We define a theory T to be falsifiably complete with respect to T' if every structure in T' which is not in T contains a finite data set which falsifies T, in the sense that no structure in T contains the same data set. To see why this definition makes sense, theory T predicts that any structure outside of T makes untrue claims. If this prediction can be empirically falsified, it means that any structure in T' but not in T should generate some finite data set which is inconsistent with T. This is precisely what our definition requires.

We show that a theory T is falsifiably complete with respect to T' if and only if it has a very specific type of axiomatization. Namely, all axioms should be universal negations of conjunctions of atomic formulas–we call this an axiomatization by UNCAF formulas. While this sounds complicated, it actually corresponds very closely with our intuition. A universal theory is, as it sounds, a theory which postulates a given relationship hold universally; that is, it is an axiom formulated with a collection of $\forall$ symbols coming at the beginning of the formula.

The proof that all falsifiably complete theories can be axiomatized by UNCAF formulas can be naively explained as follows. Each data set that is ruled out by our theory can be used to specify an axiom. This axiom states roughly that the observed data set should not occur. As each data set consists of a number of statements about the primitive elements of our language, this axiom can be written in an UNCAF form. That is, it rules out the conjunction of all statements true in the data set holding simultaneously;

and as each data set is finite, it constitutes a meaningful sentence. While this intuition is roughly correct, the proof is slightly subtle and relies on the formal definition of a theory.

Our theorem is related to the standard intuition on falsifiability; but once it is formalized, falsifiability raises some subtle issues. Popper (1959) claimed that all scientific theories must be universal, as only universal theories could be falsifiable. Popper (famously) compared two theories; the theory "all swans are white" is universal and falsifiable, while the theory "there is a non-white swan" is existential and not falsifiable. Most practitioners identify falsifiability with universality.

In fact, all falsifiably complete theories are universal, but the converse is not true. Sometimes our language lacks the expressive power to falsify our hypotheses. In other words, our data is not of the type necessary to falsify the theory. This is why we also require that the axioms be negations of conjunctions of atomic formulas (UNCAF). To understand this, suppose our theory only allows us to observe weak preference, but does not allow us to observe absence of weak preference. Consider an axiom which is simply a negation of a single atomic formula. Such an axiom might express $\forall x \forall y, \neg x \, R \, y$. The atomic formula here is $x \, R \, y$: note that the atomic formula comes in exactly the form our data comes in. Consequently, we falsify the axiom whenever we observe a pair being compared. By contrast, if our language only allows us to observe R, then we could never falsify the universal axiom: $\forall x \forall y, x \, R \, y$, as we could never observe when a pair is not compared.

1.2 Falsifiable closures

Our results on falsifiable completeness are important as they help us to characterize the falsifiable closure of a given theory. The falsifiable closure of T with respect to T' is defined to be the largest subtheory of T' with respect to which T is not falsifiable. We show that, in fact, this is a topological closure operation: it can be equivalently defined as the smallest falsifiably complete subtheory of T' containing T. So, if a theory is falsifiably complete, it is its own falsifiable closure. The relevance of this concept is that it provides the entire empirical content of a given theory. Therefore, any two theories with identical falsifiable closures are empirically indistinguishable. Hence, the theory of utility maximization and the theory of weak order are not empirically distinguishable. Other such results exist in the literature: Afriat (1967) showed that, in the theory of locally non-satiated preference maximization on budget sets, the theory of concave, monotonic utility representation of preference is empirically vacuous (see Section 6.2).

Because the falsifiable closure is a topological closure, we know that the falsifiably complete theories are closed under arbitrary intersection. Thus, for an axiomatic theory, if each axiom generates a falsifiably complete theory, then the theory characterized by all the axioms is itself falsifiably complete. We present examples to show that a collection of axioms may have more empirical content taken jointly than the intersection of their individual empirical contents, but this can only happen when falsifiable completeness fails.

1.3 Previous literature

We are not the first to formally discuss notions of falsifiability and empirical content in an abstract sense. Results exist in the mathematical psychology literature, as well as among philosophers. Adams, Fagot, and Robinson (1970) seems to be the first work discussing empirical content in a formal sense (see also Pfanzagl, Baumann, and Huber (1971) and Adams (1992)). This work defines two theories to be empirically equivalent if the set of all formulas (of a certain type) consistent with one theory is equivalent to the set of all formulas (of a certain type) consistent with the other. Just as in our work, the notion of empirical equivalence necessarily depends on what is allowed as data. The distinction is that these works do not provide a general characterization of the axiomatic structure of empirical content, but rather focus on characterizing the empirical content of specific theories. Pfanzagl, Baumann, and Huber (1971) (p. 106-119) for example, simply define testable formulas to be exactly the universal formulas.

Simon and Groen (1973) present a formal study of the testable implications of scientific theories. The focus in their work is when a theory that involves theoretical terms can be reduced to statements about observables by a process known as a Ramsey elimination. Apart from the questions that they investigate, the main difference from our work lies in their definition of data. They consider substructures (in the sense of mathematical logic) to be data. Our notion of data, on the other hand, is broader. The notion of substructure does not allow for "partial" observation, whereas our notion does. For example, given revealed preference observations of the type x is revealed preferred to y and y is revealed preferred to z, Simon and Groen would not allow the data to be silent about the relation x and y stands to each other. Our notion of data allows for such partial observation, and we believe this aspect is crucial. We discuss the alternative definition in Section 5.1, and argue that it is inadequate as a notion of data in economics.

Finally, some of our formal arguments are close to results by Tarski (1954). Tarski's main results deal with languages involving no constant or function symbols. In such a framework, he characterizes those theories that have a universal axiomatization. As we demonstrate below, the issue of universal axiomatization is related to falsification, but Tarski never explored this aspect of the results. In all, our results are hardly novel contributions to Mathematical Logic or Model Theory. Rather, we have formalized some questions that economists in particular care about, and obtained a characterization of the empirical content of a theory.

The presentation of the paper is as follows. Section 2 discusses our general notion of theory, building from concepts in model theory. Section 3 discusses our semantic notions of data, falsifiability, falsifiable completeness, and falsifiable closure. Section 4 contains our main results: syntactic characterizations of the notions presented in Section 3. The culmination of this section is Section 4.2, where we present our general results relating to relative notions of falsifiability. In Section 5, we present some related works involving Tarski. Section 6 is devoted to two applications: one involves unknown results on multiple selves; the other is a presentation of Afriat's theorem in our context. Section 7 discusses the relation of our work to the work of Simon and Groen. Lastly, Section 8 concludes. Appendix A shows how our results on falsifiability can be presented by the dual notion of verifiability, and Appendix B discusses the basic notions from mathematical logic and model theory which are required to understand our paper.

2 Theories and structures

We use standard notions from mathematical logic and model theory. To make our paper self-contained, we have included an appendix with the relevant definitions: see Appendix B. The definitions are taken quite literally from Marker (2002). At the very least, the reader should be familiar with the notions of language, structure, truth, and isomorphism of structures.

The language we choose should correspond to those objects which we believe to be observable as data in our theory. There are important and subtle issues involved in the formulation of a language. For example, for studying the basic theory of rational choice, we want a language that–at a minimum–allows us to express the observation "x is preferred to y." Thus we need a language which includes a binary relation symbol intended to represent (revealed) preference. Now, if we can observe the *absence* of preference, "x is not preferred to y," we need to include a separate relation symbol corresponding to the absence of preference. This is an important point because the absence of preference does not need to follow from the absence of an observed preference. To incorporate the observation of absence of preference, we need to incorporate this extra relation symbol. Our notion of data set (below) allows us to distinguish between the absence of observation and the observation of absence; the distinction turns out to be important.

1 Remark. We use the term 'class' for a collection that can be described by some formula in the language of set theory, but which may be 'too large' to be a set. Thus we can talk about the 'class of all sets' and 'the class of all structures of a language $\mathcal{L}$', even though these classes are not themselves sets. For a formal treatment, see Levy (2002).

2 Definition. Let $\mathcal{L}$ be a language. A *theory* T over $\mathcal{L}$ is a class of structures that is closed under isomorphism. Elements of T are called *models* of T.

3 Example. Consider the language $\mathcal{L} = \langle R, \tilde{R} \rangle$ with two binary relations:

- R, which is intended to express weak preference,
- and $\tilde{R}$, which is intended to express absence of weak preference.

A structure of L is a triple $\mathcal{M} = (M, R^{\mathcal{M}}, \tilde{R}^{\mathcal{M}})$, where M is a set, and $R^{\mathcal{M}}$ and $\tilde{R}^{\mathcal{M}}$ are binary relations on M.

The *theory of rationality* is the theory of weak-order maximization, denoted by T_{wo}. This is specified as the class of all structures $(M, R^{\mathcal{M}}, \tilde{R}^{\mathcal{M}})$ for which $R^{\mathcal{M}}$ is complete and transitive, and for all $x, y \in M$, $x \, \tilde{R}^{\mathcal{M}} \, y$ if and only if $x \, R^{\mathcal{M}} \, y$ is false. That is, $R^{\mathcal{M}}$ expresses weak preference, while $\tilde{R}^{\mathcal{M}}$ expresses the absence of weak preference.

We can write this more carefully as follows: T_{wo} is the class of all $\mathcal{L}$-structures for which the following axioms are true:

1. $\forall x \forall y, (x \, R^{\mathcal{M}} \, y) \vee (x \, \tilde{R}^{\mathcal{M}} \, y)$
2. $\forall x \forall y, \neg[(x \, R^{\mathcal{M}} \, y) \wedge (x \, \tilde{R}^{\mathcal{M}} \, y)]$
3. $\forall x \forall y \forall z, \neg[(x \, R \, y) \wedge (y \, R \, z) \wedge (x \, \tilde{R} \, z)]$
4. $\forall x \forall y, \neg[(x \, \tilde{R} \, y) \wedge (y \, \tilde{R} \, x)]$.

The first axiom expresses that there must be either preference or absence of preference between all pairs. The second axiom expresses consistency between preference and absence of preference: if there is a preference between x and y, there cannot be absence of a preference. The third formalizes transitivity, and the last formalizes completeness.

For future reference, we denote the class of all structures for which axioms 2,3, and 4 are true by T_w.

We distinguish T_{wo} from the theory of utility maximization, which is the class of $\mathcal{L}$-structures T_u for which there exists a real-valued function $u : M \to \mathbf{R}$ such that $x\,R\,y \leftrightarrow u(x) \geq u(y)$ and $x\,\tilde{R}\,y \leftrightarrow u(x) < u(y)$.

Finally, we can define the "vacuous theory" T_v of all the structures of L. Note that $T_u \subseteq T_{wo} \subseteq T_w \subseteq T_v$. So we can express that one theory is more restrictive than another by set containment.

4 Remark. Marker and other model theory textbooks only study first-order theories (See Definition 17 below). In our definition of theory we follow Tarski (1954).

3 Falsifiable Closure: Semantics

5 Definition. Let $\mathcal{L}$ be a language. A *data set* $\mathcal{D}$ over $\mathcal{L}$ is given by:

1. A non-empty set D (the domain of $\mathcal{D}$)
2. An n-ary relation $P^{\mathcal{D}}$ over D for every n-ary relation symbol P of $\mathcal{L}$
3. A function $f^{\mathcal{D}} : \mathrm{Dom}(f^{\mathcal{D}}) \subseteq D^n \to D$ for every n-ary function symbol f of $\mathcal{L}$.
4. A set $C(\mathcal{D})$ of constant symbols of $\mathcal{L}$ and an element $c^{\mathcal{D}} \in D$ for every $c \in C(D)$.

A data set $\mathcal{D}$ is *finite* if the domain D and the sets $\{P|P^{\mathcal{D}} \neq \emptyset\}$, $\{f|\mathrm{Dom}(f^{\mathcal{D}}) \neq \emptyset\}$, and $C(\mathcal{D})$ of, respectively, relation symbols, function symbols and constant symbols that *appear in* $\mathcal{D}$ are finite.

There are some subtle issues in the definition of data set. In particular, as we explain in detail in Section 5.1, a data set does not impose that one observe all the theoretically possible relations among objects in the data set. This imposition would result in a rather unrealistic notion of data set, and our definition avoids it. We model data sets in this way in order to capture the idea of partial observability.

6 Definition. Let $\mathcal{L}$ be a language. A structure $\mathcal{M}$ of $\mathcal{L}$ *contains* a data set $\mathcal{D}$, denoted $\mathcal{D} \subseteq \mathcal{M}$ if the following conditions are satisfied:

1. $D \subseteq M$, where D and M are the domains of $\mathcal{D}$ and $\mathcal{M}$.
2. $P^{\mathcal{D}} \subseteq P^{\mathcal{M}}$ for every relation symbol P
3. $f^{\mathcal{D}}$ is the restriction of $f^{\mathcal{M}}$ to $\mathrm{Dom}(f^{\mathcal{D}})$ for every function symbol f.
4. $c^{\mathcal{D}} = c^{\mathcal{M}}$ for every constant symbol $c \in C(\mathcal{D})$.

Observe that we do not require $P^{\mathcal{D}}$ to be the restriction of $P^{\mathcal{M}}$ to D (and similarly for functions). Consider the language in Example 3, and the structure $\mathcal{M} = (\mathbf{R}, \geq, <)$ of T_{wo}, where $\geq$ is the usual order on $\mathbf{R}$. Then the data set $\mathcal{D}$ with domain $\{1, 2, 3\}$ and the binary relation $R^{\mathcal{D}} = \{(2, 1)\}$, is contained in $\mathcal{M}$.
7 Definition. Let $\mathcal{L}$ be a language.

1. A data set $\mathcal{D}$ *falsifies* a theory T if no model of T contains $\mathcal{D}$.
2. Let $\mathcal{M}$ be a structure. A theory T is *falsifiable at* $\mathcal{M}$ if $\mathcal{M}$ contains a data set that falsifies T.

A theory T is falsified at a structure $\mathcal{M}$ if some claim that T makes is is incompatible with data that could be observed if $\mathcal{M}$ was the structure that represents the real world.

The following lemmas establish some simple properties which are useful later.
8 Lemma. *If $T_1 \subseteq T_2$ are theories and T_2 is falsifiable at a structure $\mathcal{M}$ then T_1 is also falsifiable at $\mathcal{M}$.*
9 Lemma. *If T_1, T_2 are theories that are falsifiable at a structure $\mathcal{M}$ then $T_1 \cup T_2$ is falsifiable at $\mathcal{M}$.*
10 Lemma. *If a theory T is falsifiable at a structure $\mathcal{M}$ then T is falsifiable at every isomorphic copy $\mathcal{M}'$ of $\mathcal{M}$.*

Proof of Lemma 8. Let $\mathcal{D} \subseteq \mathcal{M}$ be a finite data set that falsifies T_2. Then $\mathcal{D}$ falsifies T_1. □

Proof of Lemma 9. Let $\mathcal{D}_1$ and $\mathcal{D}_2$ be finite data sets that are contained in $\mathcal{M}$ and falsify T_1 and T_2 respectively. Let $\mathcal{D}_1 \cup \mathcal{D}_2$ be the data set with domain $D_1 \cup D_2$ and such that $p^{\mathcal{D}_1 \cup \mathcal{D}_2} = p^{\mathcal{D}_1} \cup p^{\mathcal{D}_2}$ for every relation symbol p, $f^{\mathcal{D}_1 \cup \mathcal{D}_2} = f^{\mathcal{D}_1} \cup f^{\mathcal{D}_2}$ for every function symbol f and $C(\mathcal{D}_1 \cup \mathcal{D}_2) = C(\mathcal{D}_1) \cup C(\mathcal{D}_2)$. Note that $f^{\mathcal{D}_1} \cup f^{\mathcal{D}_2}$ defines a function because $\mathcal{D}_1$ and $\mathcal{D}_2$ are contained in $\mathcal{M}$. Then $\mathcal{D}_1 \cup \mathcal{D}_2$ falsifies $T_1 \cup T_2$. □

Proof of Lemma 10. Let $\eta : \mathcal{M}' \to \mathcal{M}$ be an isomorphism, and let $\mathcal{D} \subseteq \mathcal{M}$ be a finite data set with domain D that falsifies T. Let $\mathcal{D}' \subseteq \mathcal{M}'$ be the data set with domain $D' = \eta^{-1}(D)$, and such that the relations and functions of $\mathcal{D}'$ are the pullbacks by η of the corresponding relations and functions of $\mathcal{D}$, $C(\mathcal{D}') = C(\mathcal{D})$ and $c^{\mathcal{D}'} = \eta^{-1}\left(c^{\mathcal{D}}\right)$ for every $c \in C(\mathcal{D})$. Then it follows from the fact that T is closed under isomorphisms that $\mathcal{D}'$ falsifies T. □

11 Definition. A theory T is *falsifiable* if there exists some data set that falsifies T.

A theory T is falsifiable if T makes at least one claim that can be demonstrated to be false. Consider Example 3. The theory T_u of utility maximization is falsifiable: the data set $\mathcal{D} = (D, R^{\mathcal{D}}, \tilde{R}^{\mathcal{D}})$ with domain $D = \{a, b\}$ and where $R^{\mathcal{D}} = \emptyset$ and $\tilde{R}^{\mathcal{D}} = \{(a, b), (b, a)\}$ falsifies T_u.

On the other hand, while T_u is falsifiable, not *all* its claims are falsifiable. For an example, consider the structure $\mathcal{M}_{lex} = (\mathbf{R}^2_+, \geq_{lex}, <_{lex})$, where $\geq_{lex}$ is the lexicographic order on $\mathbf{R}^2_+$. It is well-known that $\mathcal{M}_{lex} \notin T_u$, but no finite data set in $\mathcal{M}_{lex}$ falsifies T_u.

Thus, we may be interested in theories all of whose claims are falsifiable, and more importantly, in the empirical content of a theory such as T_u. These observations motivate the following definitions.

12 Definition. A theory T is *falsifiably complete* if T is falsifiable at every structure which is not a model of T.

13 Definition. Let T be a theory. The *falsifiable closure* of T, denoted $\mathrm{fc}(T)$ is the class of all structures $\mathcal{M}$ such that T is not falsifiable at $\mathcal{M}$.

From Lemma 10 it follows that $\mathrm{fc}(T)$ is a theory (*i.e.* closed under isomorphism). The theory $\mathrm{fc}(T)$ captures our idea of empirical content. In particular, T is falsifiably complete if and only if $\mathrm{fc}(T) = T$.

14 Example. Consider again Example 3. Then $\mathrm{fc}(T_u) = \mathrm{fc}(T_{wo}) = T_w$. Thus, the theory of utility maximization and the theory of preference maximization are empirically indistinguishable. In addition, the empirical content of T_u and T_{wo} is, in a sense, contained in axioms 2-4 of Example 3. Axiom 1 expresses a non-falsifiable property, and the additional hypotheses implicit in T_u are also non-testable.

15 Lemma. *If a theory T is falsifiable at a structure $\mathcal{M}$ then $fc(T)$ is also falsifiable at $\mathcal{M}$.*

Proof. Let $\mathcal{D}$ be a finite data set that is contained in $\mathcal{M}$ and falsifies T. By Definition 13 no model of $\text{fc}(T)$ contains $\mathcal{D}$ (since $\mathcal{D}$ falsifies T). By Definition 7 this means that $\mathcal{D}$ falsifies $\text{fc}(T)$. Since $\mathcal{M}$ contains $\mathcal{D}$ it follows that $\text{fc}(T)$ is falsifiable at $\mathcal{M}$. □

The following proposition says that the operator $T \mapsto \text{fc}(T)$ over theories T has the properties of a topological closure. The theory $\varnothing$ is the theory which contains no structures.

16 Proposition. *The falsifiable closure has the following properties.*

Extensiveness $T \subseteq fc(T)$ *for every theory* T*.*

Idempotence $fc(fc(T)) = fc(T)$ *for every theory* T*.*

Preservation of Nullary Union $fc(\varnothing) = \varnothing$*.*

Preservation of Binary Union $fc(T_1 \cup T_2) = fc(T_1) \cup fc(T_2)$ *for all theories* T_1, T_2*.*

Proof. Extensiveness follows from the fact that T is not falsifiable at its own models. Idempotence from Lemma 15: If $\mathcal{M} \notin \text{fc}(T)$ then T is falsifiable at $\mathcal{M}$ and therefore $\text{fc}(T)$ is falsifiable at $\mathcal{M}$, *i.e.* $\mathcal{M} \notin \text{fc}(\text{fc}(T))$. Preservation of nullary union follows as every model contains a data set falsifying $\varnothing$. Preservation of binary union follows from Lemma 9. □

4 Syntax

We now formalize the assertions that can be expressed using the language $\mathcal{L}$ to describe properties of $\mathcal{L}$-structures. This follows the details in Appendix B. The only departure we make from classical model theory is the inclusion of a symbol '$\neq$' in our meta-language, which is always interpreted in the "correct" way. Hence, expressions in our language are strings of symbols built from the symbols of $\mathcal{L}$, variable symbols $v_1, v_2, \ldots$, the equality and inequality symbols $=, \neq$, Boolean connectives $\neg, \vee, \wedge$, quantifiers $\exists, \forall$ and parentheses $($, $)$. As we allow the symbol $\neq$ to appear in our sentences, we need to make small changes in our definitions of term, formula, sentence, and truth. The changes necessary should be obvious to those familiar with mathematical logic; again, details are presented in Appendix B.

17 Definition. For a set Γ of sentences of $\mathcal{L}$, let $\mathcal{T}(\Gamma)$ be the theory of all structures $\mathcal{M}$ of $\mathcal{L}$ such that all the formulas in Γ are true in $\mathcal{M}$. Theories of the form $\mathcal{T}(\Gamma)$ for some set Γ of formulas are called *first-order theories*. We also say that Γ *axiomatizes* $\mathcal{T}(\Gamma)$.

18 Example. In Example 3, the theory T_{wo} is a first order theory. The theory T_u is not a first order theory. That T_u has no first order axiomatization may not be immediately obvious, but follows from classical results in model theory.

19 Definition. Let $\mathcal{L}$ be a language. A *universal negation of a conjunction of atomic formulas (UNCAF)* sentence of $\mathcal{L}$ is a sentence of the form

$$\forall v_1 \forall v_2 \ldots \forall v_n \neg (\phi_1 \wedge \phi_2 \cdots \wedge \phi_m)$$

where $\phi_1, \phi_2, \ldots, \phi_m$ are atomic formulas with variables $v_1, \ldots, v_n$.

The following result provides the syntactic characterization of the semantic concept of falsifiable completeness. Falsifiably complete theories are exactly those which have an UNCAF axiomatization. This is our main result.

20 Theorem. *A theory T is falsifiably complete if and only if it admits an UNCAF axiomatization.*

The following corollary is an immediate consequence of Theorem 20 and Definition 12. It will be of interest to us later, in comparing our work with that of Tarski (1954).

21 Corollary. *Let $\mathcal{L}$ be a language and T a theory over $\mathcal{L}$. Then T admits an axiomatization by UNCAF sentences if and only if the following condition is satisfied: For every structure $\mathcal{M}$, if every finite sub data-set of $\mathcal{M}$ is contained in some model of T then $\mathcal{M}$ is a model of T.*

The following corollary deals with finite axiomatizations. One should not necessarily expect a theory to have a finite axiomatization, as it is equivalent to a uniform bound on the size of a falsifying data set. For example in classical demand theory (Section 6.2), the theory axiomatized by the weak axiom of revealed preference can always be falsified by two observations; the strong axiom, on the other hand, is an infinite collection of axioms, and there is no bound on a falsifying data set. We took the main idea in Corollary 22 from Vaught (1954); the proof follows from the proof of Theorem 20 and it is omitted.

22 Corollary. *Let $\mathcal{L}$ be a language with finitely many symbols, and T a theory over $\mathcal{L}$. Then T admits an axiomatization by finitely many UNCAF sentences if and only if the following condition is satisfied: There is an n such that, for every structure $\mathcal{M}$, if every finite sub data-set of $\mathcal{M}$, whose domain has at most n elements, is contained in some model of T then $\mathcal{M}$ is a model of T.*

For a theory T denote by uncaf(T) the set of UNCAF formulas that are true in all models of T.

Theorem 20 is an immediate consequence of Proposition 23

23 Proposition. *For every theory T one has $fc(T) = \mathcal{T}\,(uncaf(T))$.*

Let $\mathcal{L}$ be a language and $\mathcal{D}$ a finite data set. For every $d \in D \setminus C(\mathcal{D})$ let v_d be a variable, and let z_d for every $d \in D$ be the term given by $z_d = c$ if $d = c^{\mathcal{D}}$ for some $c \in C(\mathcal{D})$ and $z_d = v_d$ if $d \in D \setminus C(\mathcal{D})$. Let $\phi_{\mathcal{D}}$ be the following UNCAF formula of $\mathcal{L}$:

$$\begin{aligned} &\phi_{\mathcal{D}} = \forall \bar{v} \neg \bar{\phi}_{\mathcal{D}}(\bar{v}), \text{ where} \\ &\bar{\phi}_{\mathcal{D}}(\bar{v}) = \left(\bigwedge (z_d \neq z_{d'}) \bigwedge P\left(z_{d_1}, \ldots, z_{d_n}\right) \wedge \bigwedge f\left(z_{d_1}, \ldots, z_{d_n}\right) = z_{f^{\mathcal{D}}(d_1,\ldots,d_n)} \right), \end{aligned} \tag{1}$$

The first conjunction ranges over all pairs $d \neq d' \in D$; the second conjunction ranges over all relation symbols P that appear in $\mathcal{D}$ and every $(d_1, \ldots, d_n) \in P^{\mathcal{D}}$; and the third conjunction ranges over all function symbols f that appear in $\mathcal{D}$ and every $(d_1, \ldots, d_n) \in \mathrm{Dom}(f^{\mathcal{D}})$.

24 Lemma. *Let $\mathcal{D}$ be a finite data set. Then $\phi_{\mathcal{D}}$ is not true in $\mathcal{M}$ if and only if $\mathcal{D}$ is contained in some isomorphic copy of $\mathcal{M}$.*

Proof of Proposition 23. We divide the proof into two steps:

Step 1: If $\mathcal{M} \in \mathcal{T}(\text{uncaf}(T))$ then $\mathcal{M} \in \text{fc}(T)$.

Let $\mathcal{D}$ be a data set that falsifies fc(T). Then from Lemma 24, and the fact that T is closed under isomorphism it follows that $\phi_{\mathcal{D}} \in \text{uncaf}(T)$. Therefore $\mathcal{M} \models \phi_{\mathcal{D}}$, as by hypothesis $\mathcal{M} \in \mathcal{T}(\text{uncaf}(T))$. By Lemma 24 again it follows that $\mathcal{M}$ does not contain $\mathcal{D}$. Therefore $\mathcal{M}$ does not contain any data set that falsifies $\mathcal{D}$, so that T is not falsifiable at $\mathcal{M}$, *i.e.* $\mathcal{M} \in \text{fc}(T)$ as desired.

Step 2: If $\mathcal{M} \notin \mathcal{T}(\text{uncaf}(T))$ then $\mathcal{M} \notin \text{fc}(T)$.

Let $\phi \in \mathcal{T}\,(\text{uncaf}(T))$ be not true in $\mathcal{M}$. Let $\bar{v} = (v_1, \ldots, v_n)$ be the variables of ϕ so that $\phi = \forall \bar{v} \neg \bar{\phi}(\bar{v}) \in \mathcal{T}\,(\text{uncaf}(T))$ where $\bar{\phi}(\bar{v})$ is a conjunction of atomic formulas.

Since ϕ is not true in $\mathcal{M}$, it follows that then $\bar{\phi}[\bar{d}]$ is true in $\mathcal{M}$ for some $\bar{d} = (d_1, \ldots, d_n)$. Let $\mathcal{D}$ be the finite data set defined as follows: The domain $D \subseteq \mathcal{M}$ of $\mathcal{D}$ is the set of all elements of the form $t[d_1, \ldots, d_k]$ where t is some term that appears in $\bar{\phi}$. For every relation symbol P,

$$P^{\mathcal{D}} = \{(t_1[d_1, \ldots, d_k], \ldots, t_n[d_1, \ldots, d_k]) \,|\, P(t_1, \ldots, t_n) \text{ appears in } \bar{\phi}\}.$$

For every function symbol f,

$$\mathrm{Dom}(f^{\mathcal{D}}) = \{(t_1[d_1, \ldots, d_k], \ldots, t_n[d_1, \ldots, d_k]) \,|\, f[t_1, \ldots, t_n] \text{ appears in } \bar{\phi}\},$$

and for every $(t_1, \ldots, t_n)$ such that the atomic formula $t = f(t_1, \ldots, t_n)$ appears in $\bar{\phi}$

$$f^{\mathcal{D}}\,(t_1[d_1, \ldots, d_k], \ldots, t_n[d_1, \ldots, d_k]) = t[d_1, \ldots, d_k].$$

If there are two different atomic formulas that appear in $\bar{\phi}$ with the same arguments of f then we choose one of them arbitrarily to define the corresponding value of $f^{\mathcal{D}}$.

Then $\mathcal{D}$ is a data set that is contained in $\mathcal{M}$ and $\bar{\phi}[d_1, \ldots, d_k]$ is true in every structure that contains $\mathcal{D}$, and, in particular, ϕ is not true in any structure that contains $\mathcal{D}$. But ϕ is is true in every model of T, and therefore $\mathcal{D}$ falsifies T. Thus, we proved that $\mathcal{M}$ contains the data set $\mathcal{D}$ that falsifies T and therefore $\mathcal{M} \notin \mathrm{fc}(T)$. □

Proof of Lemma 24. If a structure $\mathcal{M}$ contains $\mathcal{D}$ then substituting d for v_d we get that $\bar{\phi}_{\mathcal{D}}[\bar{d}]$ is false in $\mathcal{M}$ and therefore $\phi_{\mathcal{D}}$ is not true in $\mathcal{M}$. Since truth is preserved under isomorphism, it follows that if an isomorphic copy of $\mathcal{M}$ contains $\mathcal{D}$ then $\phi_{\mathcal{D}}$ is not true in $\mathcal{M}$.

Assume now that $\mathcal{M}$ is a structure of $\mathcal{L}$ such that $\phi_{\mathcal{D}}$ is not true in $\mathcal{M}$, and assume without loss of generality that the domains M and D of $\mathcal{M}$ and $\mathcal{D}$ are disjoint (otherwise replace $\mathcal{M}$ with an isomorphic structure). Let $\bar{m} = (m_d)_{d \in D}$ be elements of $\mathcal{M}$ such that $\bar{\phi}_{\mathcal{D}}[\bar{m}]$ is false in $\mathcal{M}$. Consider the isomorphic structure of $\mathcal{M}'$ which is obtained by replacing every element m_d with d. Then $\bar{\phi}_{\mathcal{D}}[\bar{d}]$ is false in $\mathcal{M}'$. It follows that all the corresponding substitutions of $\bar{d}$ in the atomic formulas in the conjunctions that makes up $\phi_{\mathcal{D}}$ in (1) are true. In particular, $(d_1, \ldots, d_n) \in P_{\mathcal{M}'}$ for every relation symbol P that appears in $\mathcal{D}$ and every $(d_1, \ldots, d_n) \in P^{\mathcal{D}}$. Thus, $P^{\mathcal{D}} \subseteq P_{\mathcal{M}'}$ for every relation symbol P that appears in $\mathcal{D}$, and so property (2) in Definition 6 is satisfied. The other properties are proved by similar argument. Therefore $\mathcal{M}'$ is an isomorphic copy of $\mathcal{M}$ that contains

$\mathcal{D}$. □

4.1 Joint hypotheses

We present a trivial example establishing that the falsifiable closure operator does not commute with respect to intersection. While the falsifiable closure of two falsifiably complete theories is the intersection of the closures, this is not true of theories that are not falsifiably complete.

25 Example. Let the language $L = \langle R, S \rangle$ involve two unary relations. T' is the vacuous theory of all structures with two unary relations. T_1 is the theory axiomatized by $\forall x, R(x)$. T_2 is the theory axiomatized by $\forall x, R(x) \to \neg S(x)$. Note that the falsifiable closure of T_1 is T', while the falsifiable closure of T_2 is T_2 itself. Consequently, the intersection of the falsifiable closures is T_2.

However, the UNCAF axiom $\forall x, \neg S(x)$ is true in $T_1 \cap T_2$, while it is not true in either T_1 or T_2. Consequently the falsifiable closure of $T_1 \cap T_2$ is a proper subtheory of the intersection of the individual falsifiable closures.

The example is trivial, but captures the essence of a familiar problem. It is possible that two theories imposed jointly imply stronger hypotheses than just those which follow logically from each of the two theories. Our results imply that this only happens for theories which are not falsifiably complete.

4.2 Relative notions

It is often useful to have a relative notion of falsifiability. In some cases, there is a theory which we postulate to be a "base" theory, and we want to test some additional hypothesis (a stronger theory). For example, consider the theories in Example 3. We may ask about additional empirical content in the T_u, relative to T_{wo}; and conclude that the hypotheses that T_u adds to T_{wo} have no additional empirical content.

The theories we have been describing up until now must be necessarily completely specified, and everything that these theories postulate must be open to testing–including the primitives. Our results do not require such a detailed description.

To take a trivial example, we may know that there are at least three alternatives over which an agent forms a preference. We could formalize this by ensuring that all structures in our theory have universes with at most three elements. It turns out that so long as our theory is not vacuous, this theory could never be falsifiably complete. The reason is that, if we are given any model $\mathcal{M}$ of our theory, and consider a substructure $\mathcal{M}^* \subseteq \mathcal{M}$ of this theory with a universe containing only two elements, then $\mathcal{M}^*$ is clearly not a model of our theory. But our theory is also not falsifiable at $\mathcal{M}^*$, as $\mathcal{M}^* \subseteq \mathcal{M}$. This is only a trivial example, of course, but it illustrates the need to allow for some hypotheses to be taken as "given."[5]

To discuss relative notions of falsifiability, in this section we fix two theories $T \subseteq T'$. We assume that T' is a "base", or known, theory. We say that T is *falsifiable with respect to* T' if T is falsifiable at some model of T'. Thus a theory T is falsifiable with respect to a weaker theory T' if some claim that T makes in addition to T' is incompatible with data that could be observed if T' were true. T is *falsifiably complete with respect to* T' if T is falsifiable at every model of T' which is not a model of T. The *falsifiable closure of* T *in* T', denoted $\mathrm{fc}_{T'}(T)$, is given by $\mathrm{fc}_{T'}(T) = T' \cap \mathrm{fc}(T)$, the class of all models $\mathcal{M}$ of T' such that T is not falsifiable at $\mathcal{M}$. Note that T is falsifiably complete with respect to T' if and only if $\mathrm{fc}_{T'}(T) = T$. We have the following theorem:

26 Theorem. *Suppose $T \subseteq T'$. Then T is falsifiably complete with respect to T' if and only if there exists a set Σ of UNCAF sentences of $\mathcal{L}$ such that $T = T' \cap \mathcal{T}(\Sigma)$.*

Proof of Theorem 26. If T is falsifiably complete with respect to T', then by Proposition 23

$$T = \mathrm{fc}_{T'}(T) = T' \cap \mathrm{fc}(T) = T' \cap \mathcal{T}(\Sigma),$$

where $\Sigma = \mathrm{uncaf}(T)$.

Assume now that $T = T' \cap \mathcal{T}(\Sigma)$ for some set Σ of UNCAF sentences. In particular, every sentence in Σ is true in every model of T and therefore $\Sigma \subseteq \mathrm{uncaf}(T)$. It follows that

$$T' \cap \mathrm{fc}(T) = T' \cap \mathcal{T}\left(\mathrm{uncaf}(T)\right) \subseteq T' \cap \mathcal{T}(\Sigma) = T,$$

where the first equality follows from Proposition 23 and the inclusion from the fact that $\Sigma \subseteq \mathrm{uncaf}(T)$. Since in addition $T \subseteq T' \cap \mathrm{fc}(T)$, we get that $T = T' \cap \mathrm{fc}(T)$, so that T is falsifiably complete with respect to T'. □

27 Proposition. *Let $T \subseteq T'$ be theories. Then $fc_{T'}(T)$ is the smallest theory that contains T and is falsifiably complete with respect to T'.*

[5]We may also decide to take some mathematical objects as given, so that our axiomatization only needs to characterize economically meaningful hypotheses.

Proof. From the fact that fc is idempotent and monotone (Proposition 16), we conclude that

$$\mathrm{fc}_{T'}(\mathrm{fc}_{T'}(T)) = \mathrm{fc}(\mathrm{fc}(T) \cap T') \cap T' \subseteq \mathrm{fc}(\mathrm{fc}(T)) \cap T' = \mathrm{fc}(T) \cap T' = \mathrm{fc}_{T'}(T).$$

Therefore $\mathrm{fc}_{T'}(T)$ is falsifiably complete with respect to T'. Assume now that $T \subseteq \tilde{T} \subseteq T'$ and $\tilde{T}$ is falsifiably complete with respect to T'. Then

$$\mathrm{fc}_{T'}(T) \subseteq \mathrm{fc}_{T'}(\tilde{T}) = \tilde{T},$$

where the first inclusion follows from monotonicity of the closure and the fact that $T \subseteq \tilde{T}$ and the equality from the fact that $\tilde{T}$ is falsifiably complete with respect to T'. □

28 Example. Consider again the language $\mathcal{L} = \langle R, \tilde{R} \rangle$. We define the theory of orders, T_o, as the class of all structures satisfying

$$\forall x \forall y, [(x\ R^{\mathcal{M}}\ y) \leftrightarrow \neg(x\ \tilde{R}^{\mathcal{M}}\ y)].$$

Then $\mathrm{fc}_{T_o}(T_u) = \mathrm{fc}_{T_o}(T_{wo}) = T_{wo}$. That is, if we assume that every pair is either ranked or unranked (in fact, this assumption would usually be implicit), then the theory of weak order is falsifiably complete. The theory of weak order is the falsifiable closure of the theory of utility maximization. The idea that numerical representation of preference is without empirical content is well-known, but it is comforting that our formal notion coincides with our intuition in this case.

4.3 A result on axiomatizations using unobservables

Often, a theory has an axiomatization involving unobservables. Obviously, such an axiomatization cannot directly lead to empirical falsification. We find conditions under which a theoretical axiomatization can be "projected" on observables to yield a falsifiably complete theory.

Let $\mathcal{F} \subseteq \mathcal{L}$ be languages, such that $\mathcal{L}$ contains all the symbols of $\mathcal{F}$ and possibly additional relation symbols. The idea is that the additional symbols in $\mathcal{L}$ are meant to signify theoretical and unobservable terms. For every $\mathcal{L}$-structure $\mathcal{M}$, we denote by $F(\mathcal{M})$ the $\mathcal{F}$-structure induced from $\mathcal{M}$ by forgetting the relations that corresponds relation symbols not in $\mathcal{F}$. For every $\mathcal{L}$-theory T we denote by $F(T)$ the theory of all structures of the form $F(\mathcal{M})$ for some model $\mathcal{M}$ of T.

29 Proposition. *If T is a falsifiably complete $\mathcal{L}$-theory then $F(T)$ is a falsifiably complete $\mathcal{F}$-theory.*

So a theory that is falsifiably complete when we say that theoretical objects are observable is automatically falsifiably complete in the correct observable form—as long as the observable structures are obtained by "projection" from unobservables as in the proposition.

5 Relation to Tarski

5.1 Data sets vs. substructures

Our notion of data sets have an important feature. One may only be able to observe some relations among the data, not all of them. For example, for data on revealed preferences, if one observes that x is revealed preferred to y, and that y is revealed preferred to z, one may not know (not observe) the direction of revealed preference between x and z. Our notion of data sets accommodates this feature of real-world data sets. The competing notion of substructures as data sets (see the discussion in Section 7) does not.

5.2 Tarski's result on relational systems

An UNCAF sentence is a special case of a *universal sentence*, i.e. a sentence of the form

$$\forall v_1 \dots v_n \phi(v_1, \dots v_n),$$

where ϕ is quantifier-free formula. A theory T admits a universal axiomatization if $T = \mathcal{T}(\Sigma)$ for some set Σ of universal sentences.

Tarski (1954) proved the following theorem:

30 Theorem. *Let $\mathcal{L}$ be a language without constants and function symbols and let T be a theory over $\mathcal{L}$. Then T admits axiomatization by universal sentences if and only if the following conditions are satisfied:*

1. *T is closed under substructures.*
2. *For every structure $\mathcal{M}$, if every finite substructure of $\mathcal{M}$ is a model of T, then $\mathcal{M}$ is a model of T.*

The similarity of our condition in Corollary 21 and Tarski's second condition is clear: In our framework data sets replace substructures. Indeed; the reason we are able to prove a theorem axiomatizing theories with function symbols whereas Tarski could not is that the notion of data set allows a function to be defined on a subdomain of the universe under consideration. In general; however, if we consider a function restricted to an arbitrary subset of a universe, the function may not take values in that subset, and hence the resulting object will not be a substructure. In a sense, the distinction between functions and relations in mathematical logic is made because of the way these objects relate across structures: in our context, they can be considered the same type of object (any function is a relation).

We now turn to formalize the relationship between the syntactic notions of UNCAF and universal axiomatization.

Let us say that a language $\mathcal{L}$ *supports negation of relations* if its relation symbols are divided into pairs $(P, \tilde{P})$. The idea is that $\tilde{P}$ should represent the relation 'P does not hold'. If $\mathcal{L}$ supports negation of relations, we denote by $\Lambda_{\mathcal{L}}$ the set of sentences of the form

$$\forall v_1 \dots \forall v_n \neg P(v_1, \dots, v_n) \leftrightarrow \tilde{P}(v_1, \dots, v_n)$$

for all n-ary relation symbols p in the language. We say that a theory T *respects negation of relations* if $T \subseteq \mathcal{T}(\Lambda_{\mathcal{L}})$, so that $\tilde{P}$ is interpreted as 'P does not hold' in all models of T.

31 Lemma. *Let $\mathcal{L}$ be a language that supports negation of relations. Then for every universal sentence ϕ in $\mathcal{L}$ there exist UNCAF sentences $\phi_1, \dots, \phi_n$ such that $\Lambda_{\mathcal{L}} \vdash \phi \leftrightarrow \phi_1 \wedge \dots \wedge \phi_n$.*

32 Corollary. *Let $\mathcal{L}$ be a language that supports negation of relations, and let $T \subseteq \mathcal{T}(\Lambda_{\mathcal{L}})$. Then there exists a set of universal sentences Σ such that $T = \mathcal{T}(\Lambda_{\mathcal{L}}) \cap \mathcal{T}(\Sigma)$ if and only if there exists a set of UNCAF sentences Σ' such that $T = \mathcal{T}(\Lambda_{\mathcal{L}}) \cap \mathcal{T}(\Sigma)$. T admits a universal axiomatization relative to $\mathcal{T}(\Lambda_{\mathcal{L}})$ if and only if T admits an UNCAF axiomatization relative to $\mathcal{T}(\Lambda_{\mathcal{L}})$.*

Thus, for theories that respect negation of relations our theorem and Tarski's provide the same type of axiomatization.

Proof of Lemma 31. We give a purely syntactic proof: Consider the universal sentence $\forall \bar{v} \bar{\phi}(\bar{v})$, where ϕ is quantifier free and $\bar{v}$ are the variables that appear in ϕ. Writing $\bar{\phi}$ in its conjunctive normal form, we get that ϕ is equivalent to a formula of the form

$$\forall \bar{v} \bigwedge_{i=1}^{m} \bigvee_{j=1}^{n} \phi_{i,j}$$

where each $\phi_{i,j}$ is a *literal*, i.e. an atomic formula or a negation of an atomic formula. Changing the order of the conjunction and the universal quantifier we obtain a formula of the form

$$\bigwedge_{i=1}^{m} \forall \bar{v} \bigvee_{j=1}^{n} \phi_{i,j}.$$

Using De Morgan's law and replacing each $\phi_{i,j}$ with its negation we get a formula of the form

$$\bigwedge_{i=1}^{m} \forall \bar{v} \neg \bigwedge_{j=1}^{n} \phi_{i,j}. \tag{2}$$

Finally, under $\Lambda_{\mathcal{L}}$ every literal is equivalent to an atomic formula: for every term $t_0, t_1, \ldots, t_k$, $\neg f(t_1, \ldots, t_k) = t_0$ is equivalent to $f(t_1, \ldots, t_k) \neq t_0$, and $\neg P(t_1, \ldots, t_k)$ is equivalent to $\tilde{P}(t_1, \ldots, t_k)$. Therefore we can change the formulas $\phi_{i,j}$ in (2) to atomic formulas and so we arrive at a conjunction of UNCAFs, as desired. □

In fact, for the theory of falsifiability, it is often important that our theory support negation of relations. Recall Popper's theory "all swans are white." Clearly, such a theory could never be falsified if it were impossible to observe a swan which was *not* white. The following example is our example of weak order maximization, recast in a language involving only one relation.

33 Example. Let $\mathcal{L} = \langle R \rangle$ be a language involving only one binary relation, interpreted as weak preference. Consider the theory T^*_{wo}, where $\mathcal{M} = (M, R^{\mathcal{M}}) \in T^*_{wo}$ if and only if $R^{\mathcal{M}}$ is a weak order on M. Let T^*_v denote the vacuous theory, consisting of all structures with binary relations. We claim that $\mathrm{fc}(T^*_{wo}) = T^*_v$. This means, in particular, that the theory of weak order has no empirical content unless one can reasonably observe absence of preference.

To see why this is the case, let $\mathcal{D} = (D, R^{\mathcal{D}})$ be a data set, and let $\mathcal{M} = (D, R^{\mathcal{M}})$, where $R^{\mathcal{M}}$ is the binary relation which ranks all pairs. Then $\mathcal{D} \subseteq \mathcal{M}$, and $\mathcal{M} \in T^*_{wo}$.

The result seems surprising, but it says nothing more than the well-known fact that the preference which is indifferent between all alternatives can rationalize any choices whatsoever when choices are not fully observable.

5.3 The theorem of Łoś-Tarski

Theorem (Łoś-Tarski). *A first order theory is closed under substructures if and only if it admits a universal axiomatization.*

We now turn to give an analogue of Łoś-Tarski's theorem for the case of UNCAF axiomatizations. Let $\mathcal{L}$ be a language. Let $\mathcal{M}$ and $\mathcal{N}$ be structures of $\mathcal{L}$ with domains M and N respectively. Recall that $\mathcal{M}$ is a *weak substructure* of $\mathcal{N}$ if there exists an embedding $\eta : M \to N$ such that

1. $\eta\left(f^{\mathcal{M}}(a_1, \ldots, a_n)\right) = f^{\mathcal{N}}(\eta(a_1), \ldots, \eta(a_n))$ for every n-ary function symbol f

2. $(a_1, \ldots, a_n) \in R^{\mathcal{M}}$ only if $(\eta(a_1), \ldots, \eta(a_n)) \in R^{\mathcal{N}}$ for every n-ary relation symbol R

3. $\eta(c^{\mathcal{M}}) = c^{\mathcal{N}}$ for every constant symbol c.

34 Theorem. *A first order theory is closed under weak substructures if and only if it admits an UNCAF axiomatization.*

The proof is similar to the proof of Łoś-Tarski's Theorem and is omitted.

6 Applications

6.1 Application: Multiple selves preferences

We apply our concepts to a popular model without a known axiomatization, the model of multiple selves. The purpose of this exercise is to demonstrate that the concepts we introduce are useful for studying theories which have no known axiomatizations (and hence whose empirical content is not completely understood). Models of multiple selves are motivated by empirical observations (see e.g. Ambrus and Rozen (2008), Green and Hojman (2008), Manzini and Mariotti (2007), O'Donoghue and Rabin (1999) or Fudenberg and Levine (2006)), but often they lack an axiomatization in terms of observables. Here we exhibit a broad class of such models which are falsifiably complete.

Given is a fixed and finite set of agents, the "selves." Given is also a rule for aggregating agents' preferences into a single preference. The interpretation is that an individual has conflicting preferences (perhaps different preferences for different motivations) and reconciles these preferences with a preference aggregation rule. We observe an aggregate preference (a revealed preference), and we would like to know whether it could be generated by the rule for *some* profile of agents' preferences.[6] We want to test whether or not a specific group of selves uses a particular preference aggregation rule in making decisions, only having observed the aggregate ranking. This question is the correct formulation of the standard revealed-preference exercise for the multiple selves model.

Multiple selves theories are an excellent example of how hard it can be to show falsifiability. The theories have a trivial existential (second-order) axiomatization: Given a preference aggregation rule, the theory is the collection of observables for which *there exists* preferences for individual selves generating the observable behavior. Leaving aside the second-order nature of this axiomatization, the problem with an existential axiomatization is that we cannot conclude that the theory is falsifiable. Recall the example of Popper (1959): the theory that there is a non-white swan is not testable because we would need to examine all the swans in the universe. Here, for a given observed behavior, we would need to check all possible preferences that the selves might have; for an infinite set of alternatives, this set of preferences is vast. The fact that the axiomatization is second order means we have to search over preference profiles–themselves extremely complicated objects. We present a class of aggregation rules that lead to falsifiably complete theories; theories with an UNCAF axiomatization.

We require a *finite* cardinality of agents, and any preference aggregation rule which is neutral and satisfies independence of irrelevant alternatives. We show that the theory is falsifiably complete, given that we can observe both aggregate preference and absence of aggregate preference (and that these relations behave in the proper way).

The models relate to the theory of social choice, where there have been efforts to axiomatize relations which are so rationalizable. When the society can be arbitrarily large, it is known that *any* transitive antisymmetric relation is the Pareto relation for *some* society (which may be large)–this is essentially the Szpilrajn theorem. Because of this, any complete binary relation with a transitive asymmetric part is the result of the Pareto extension rule for some society (we identify indifferent alternatives for the Pareto extension rule with unranked alternatives for the Pareto ordering–see Sen (1969)). Results for majority rule are even weaker: McGarvey (1953) showed that any complete binary relation is the majority rule relation for some society of agents (which again may

[6]In this paper, we focus on preferences which are linear orders; however the results apply more broadly.

be large). Kalai (2004) generalizes this result to a much broader class of social choice rules.

The current behavioral literature interprets the society as a group of conflicting tendencies within an individual decision maker: multiple selves. This literature attempts to understand the empirical content of such assumptions. In particular, Green and Hojman (2008) generalize McGarvey's program to choice functions. Ambrus and Rozen (2008) give sufficient conditions (stated in terms of number of "violations" of classical rationality) for a choice function to be rationalizable by conflicting selves for a fixed number of agents. DeClippel and Eliaz (2009) provide a full characterization of choice rules which can result from a specific social choice rule–the fallback solution on a *fixed pair* of agents. We only consider preference relations and not choice functions here; however, we show that the predictions of nearly every such model can be empirically falsified even in the case where we hypothesize a finite and known cardinality of "selves."[7]

There are very few results like ours, assuming a fixed and finite population of selves. Dushnik and Miller (1941) give necessary and sufficient conditions for a binary relation to be the intersection of a pair of linear orders; this can dually be seen as an axiomatization for binary relations which are the image of the Pareto extension rule for two agents. This characterization theorem both relies on existential quantification, and is not a first order characterization.[8] Dushnik and Miller (1941)'s existential axiomatization cannot be the basis for falsification.[9] Sprumont (2001) provides a similar characterization in a restricted case. Both of these results are of interest as those relations which are the intersection of a pair of linear orders are exactly those relations which can be rationalized by the Pareto-extension rule.

We work with neutral preference aggregation rules which satisfy independence of irrelevant alternatives. By working with such preference aggregation rules, we need not specify what the global set of alternatives is in advance. A set of agents N is fixed and finite. A *preference aggregation rule* is therefore defined to be a mapping carrying any set of alternatives X and any N vector of linear orders[10] (termed a *preference profile*) over those alternatives $(R^1, ..., R^n)$ to a complete binary relation over X. We write $R_{f(R^1,...,R^n)}$ for the binary relation which results (suppressing notation for dependence on X). We

[7] It is surprisingly more difficult to axiomatize such models for a fixed and known set of selves, than for an arbitrary set of agents.

[8] That is, it involves quantification over relations.

[9] In particular, the theory of Pareto relations for n agents was not known to be falsifiably complete. Our Theorem 36 demonstrates that it is.

[10] A linear order is complete, transitive, and anti-symmetric

assume the following property:

35 Definition. (Neutrality and Independence of irrelevant alternatives): For all sets X and Y, for all $x, y \in X$ and all $w, z \in Y$ and all preference profiles $(R^1, ..., R^n)$ over X and $(R'^1, ..., R'^n)$ over Y, if for all $i \in N$, $x\ R^i\ y \Leftrightarrow w\ R'^i\ z$, then $x\ R_{f(R^1,...,R^n)}\ y \Leftrightarrow w\ R_{f(R'^1,...,R'^n)}\ z$.[11]

This hypothesis embeds both the neutrality and independence of irrelevant alternatives assumptions. These assumptions seem to be the minimal assumptions needed to apply Theorem 30.

Given f, we will say that a binary relation R on a set X is *f-rationalizable* if there exists a profile of linear orders $(R_1, ..., R_n)$ for which $R = R_{f(R^1,...,R^n)}$.

Denote by $\mathcal{L} = \langle R, \tilde{R} \rangle$ the language involving two binary relations, and let $\mathcal{T}(\Lambda_{\mathcal{L}})$ be the theory of all structures satisfying the axiom

$$\forall x \forall y, xR^{\mathcal{M}}y \leftrightarrow \neg x\tilde{R}^{\mathcal{M}}y.$$

A structure is *f-rationalizable* if $R^{\mathcal{M}}$ is f-rationalizable and $xR^{\mathcal{M}}y \leftrightarrow \neg x\tilde{R}^{\mathcal{M}}y$. The class of f-rationalizable structures is denoted $\mathcal{T}_f$. Note that $\mathcal{T}_f$ is in fact a theory, as it is closed under isomorphism (this is the content of neutrality).

36 Theorem. *For every f, $\mathcal{T}_f$ is falsifiably complete with respect to $\mathcal{T}(\Lambda_{\mathcal{L}})$.*

37 Remark. If absence of ranking is unobservable, that is if we consider the language that include only a single relation symbol R, then the theory of all structures in which R is an aggregation of n linear orders is not falsifiably complete. The easiest example is when the aggregation rule is such that $xR_{f(R_1,...,R_n)}y$ for every $x, y, R_1, \ldots, R_n$. Then the theory is axiomatized by $\forall x \forall y xRy$, which is not falsifiably complete.

Proof. We first show that $\mathcal{T}_f$ has a universal axiomatization; the result then follows immediately from Corollary 32.

We use Theorem 30 to show that $\mathcal{T}_f$ is universally axiomatizable.

We must verify that T_f satisfies the following two properties:

[11] Formally, neutrality means that social rankings should be independent of the names of alternatives, and independence of irrelevant alternatives means that the social preference between a pair of alternatives should depend only on the individual preferences between that pair. We have collapsed these two hypotheses into one larger condition.

1. Closure under substructures: If $\mathcal{A} \in \mathcal{T}$, and $\mathcal{A}'$ is a substructure of $\mathcal{A}$, then $\mathcal{A}' \in \mathcal{T}$.[12]

2. Finite substructure property: If for all finite substructures $\mathcal{A}'$ of $\mathcal{A}$, $\mathcal{A}' \in \mathcal{T}$, then $\mathcal{A} \in \mathcal{T}$.[13]

The first property is obviously satisfied; it follows from the neutrality and IIA assumption.

To prove that the second is satisfied, let $\mathcal{A}$ be an arbitrary structure for the language $\mathcal{L}$, and suppose that for all finite substructures $\mathcal{A}'$ of $\mathcal{A}$, $\mathcal{A}' \in \mathcal{T}$. A structure consists of a set X and a complete binary relation $R^{\mathcal{M}}$ on X, where $\tilde{R}^{\mathcal{M}}$ is a binary relation which is the complement of $R^{\mathcal{M}}$. The assumption that for all finite substructures $\mathcal{A}'$, $\mathcal{A}' \in \mathcal{T}$ means that for all finite subsets $Y \subseteq X$, $R^{\mathcal{M}}|_Y$ is f-rationalizable. We need to show that $R^{\mathcal{M}}$ on X is also f-rationalizable.

To this end, consider $\{0, 1\}$ endowed with the discrete topology. Identify the set of binary relations on X with $\mathcal{B} = \{0, 1\}^{X \times X}$ and topologize with the product topology. Then $\mathcal{B}$ is a compact topological space. Denote the set of f-rationalizable binary relations on X by $\mathcal{B}_f$. For $x, y \in X$, let $B_{x,y} = \{B \in \mathcal{B}_f : B(x, y) = R^{\mathcal{M}}(x, y)\}$. Note that for all (x, y), $B_{x,y}$ is nonempty.[14] We now seek to show that it is a closed subset of $\mathcal{B}$. To see this, note that for all $B \in B_{x,y}$, by definition, there exists a preference profile $(R^1, ..., R^n)$ for which $B = R_{f(R^1,...,R^n)}$. For each $B \in B_{x,y}$, choose one such profile. Suppose $\{B_\lambda\}_{\lambda\in\Lambda} \subseteq B_{x,y}$ is a net converging to some $\overline{B}$. By compactness of $\mathcal{B}$, we may without loss of generality assume that $(R^1_\lambda, ..., R^n_\lambda) \to (\overline{R^1}, ..., \overline{R^n})$ (see Kelley (1955), p. 71). In particular, it is easy to verify that each $\overline{R^i}$ is a linear order (by definition of product topology convergence). Also by definition $\lim_{\lambda\in\Lambda} R_{f(R^1_\lambda,...,R^n_\lambda)} = B$. The limit can be passed through f.[15] Conclude that $R_{f(\overline{R^1},...,\overline{R^n})} = B$. Clearly, $B(x, y) = R^{\mathcal{M}}(x, y)$. Conclude that $B \in B_{x,y}$, so that $B_{x,y}$ is closed.

[12] A structure $\mathcal{A}' = (X', R')$ is a substructure of $\mathcal{A} = (X, R)$ if $X' \subseteq X$ and $R|_{X'} = R'$.

[13] A structure $\mathcal{A} = (X, R)$ is finite if X is finite.

[14] This follows from the fact that $R^{\mathcal{M}}|_{\{x,y\}}$ is f-rationalizable. This implies that there exist linear orders $(R^1, ..., R^n)$ on $\{x, y\}$ for which $R_{f(R^1,...,R^n)} = R^{\mathcal{M}}|_{\{x,y\}}$. The argument now follows from the Szpilrajn theorem, by taking appropriate extensions of R^i for all i and appealing to independence of irrelevant alternatives.

[15] To see this, note that for all $x, y \in X$ by definition of convergence, there exists $\lambda^* \in \Lambda$ for which for all $\lambda \geq \lambda^*$ and for all $i \in N$, $R^i_\lambda(x, y) = \overline{R^i}(x, y)$. Recall that for a pair $x, y \in X$ for which $x \neq y$ and a linear order R over X, $R|_{\{x,y\}}$ is determined by $R(x, y)$. As $R^i_\lambda(x, y) = \overline{R^i}(x, y)$ for all $i \in N$ and $\lambda \geq \lambda^*$, we may conclude that $R_{f(R^1_\lambda,...,R^n_\lambda)}|_{\{x,y\}} = R_{f(\overline{R^1},...,\overline{R^n})}|_{\{x,y\}}$ for all such λ. Therefore, $R_{f(R^1_\lambda,...,R^n_\lambda)} \to R_{f(\overline{R^1},...,\overline{R^n})}$.

Now, we claim that $\bigcap_{(x,y)\in X\times X} B_{x,y} \neq \varnothing$. To show this, we will show that for every finite set $Z \subseteq X \times X$, $\bigcap_{(x,y)\in Z} B_{x,y} \neq \varnothing$ and appeal to the finite intersection property. So, let $Z \subseteq X \times X$ be finite. Let $Y = Z_1 \times Z_2$, where Z_i denotes the projection of Z on the ith coordinate. Note that Y is finite; so by hypothesis, $R^{\mathcal{M}}|_Y$ is f-Pareto rationalizable. Let $(R^1, ..., R^n)$ be linear orders on Y for which $R_{f(R^1,...,R^n)} = R^{\mathcal{M}}|_Y$. Each of these can be extended to linear orders on all of X by the Szpilrajn theorem, say to R^{i*}. Then $R_{f(R^{1*},...,R^{n*})}|_Y = R^{\mathcal{M}}|_Y$ (this follows from the neutrality and independence of irrelevant alternatives hypothesis). In particular, for all $(x, y) \in Z$, $R_{f(R^{1*},...,R^{n*})}(x, y) = R^{\mathcal{M}}(x, y)$, so that $\bigcap_{(x,y)\in Z} B_{x,y} \neq \varnothing$. This verifies the finite intersection property, and as each $B_{x,y}$ is closed and $\mathcal{B}$ is compact, we conclude that $\bigcap_{(x,y)\in X\times X} B_{x,y} \neq \varnothing$. This establishes that $R^{\mathcal{M}} \in \mathcal{B}_f$.

□

The above discussion assumes that preferences are linear orders, but many of the multiple-selves papers put different restrictions on the selves' preferences. While the proof above does not directly apply, it is easy to see that the theorem is true on different domains of preference profiles: Any domain of preference profiles which is closed in the product topology as defined above will work.

6.2 Application: Afriat's theorem

Afriat's theorem (Afriat, 1967; Varian, 1982) states that consumption data are rationalizable by a monotonic, continuous, and concave utility if and only if they are rationalizable by a locally nonsatiated preference. Similarly, for demand data satisfying Walras' Law, data which are rationalizable at all are rationalizable by a monotonic, continuous, and concave utility. We shall recast his theorem, using our results, as a statement about the empirical content (the falsifiable closure) of the theory of concave utility maximization.

The language and definitions are similar to those of Example 3, but we need to make some changes to model that preferences are revealed by demand choices at competitive budgets.

Let $\Pi \subseteq \mathbf{R}^n_{++} \times \mathbf{R}_+$. A function $d : \Pi \to \mathbf{R}^n_+$ that satisfies

1. $p \cdot d(p, I) = I$, and

2. $d(p,I) = d(\lambda p, \lambda I)$ for all $\lambda > 0$ such that $(\lambda p, \lambda I) \in \Pi$

is a *demand function*.

Let $\mathcal{L}$ be a language with two binary relations, R and P. The language should also include a constant symbol for every element of $\mathbf{R}^n_+$ and $\mathbf{R}$.[16] We shall introduce three theories: the theory T' of classical demand theory, the subtheory T'_{wo} of weak-order maximization, and the subtheory T_c of concave utility maximization.

First, T' is the class of all structures isomorphic to some $\mathcal{M}$ of $\mathcal{L}$ with $M = \mathbf{R}^n_+$, all constant symbols refer to their named objects, and for which there is a demand function d and $\Pi \subseteq \mathbf{R}^N_{++} \times \mathbf{R}_+$, such that

- $(x,y) \in R$ if and only if there is $(p,I) \in \Pi$ such that $x = d(p,I)$ and $p \cdot y \leq I$;
- $(x,y) \in P$ if and only if there is $(p,I) \in \Pi$ such that $x = d(p,I)$ and $p \cdot y < I$.

Second, the theory of weak order maximization is the subtheory T'_{wo} of T' defined as structures isomorphic to some $(\mathbf{R}^n_+, R^*, P^*)$ in T' for which there is a complete, reflexive, and transitive binary relation $\succeq$ on X such that

$$(x,y) \in R^* \Rightarrow (x,y) \in \succeq$$
$$(x,y) \in P^* \Rightarrow (x,y) \in \succ .$$

The theory of concave utility maximization is the subtheory T_c of T' that is the class of all structures isomorphic to some $(\mathbf{R}^n_+, R^*, P^*)$ in T' for which there is a monotonic and concave function $u : \mathbf{R}^n_+ \to \mathbf{R}$ such that

$$(x,y) \in R^* \Rightarrow u(x) \geq u(y)$$
$$(x,y) \in P^* \Rightarrow u(x) > u(y).$$

We obtain the following expression of Afriat's (1967) theorem:

38 Theorem. *T'_{wo} is the falsifiable closure of T_c with respect to T'.*

[16] We introduce constant symbols for each element of $\mathbf{R}^n_+$ so that we do not need to worry about describing consumption space and the relation $\geq$, the function $\cdot$, etc. as part of the problem. The technique of introducing a constant to represent every element in some concrete set is very useful in a variety of contexts in which the underlying set is something whose behavior is well-understood, but whose defining symbols are not meant to be taken as data. Otherwise, we would need to take $\geq$ and the values of the function $\cdot$ as "observable data."

Proof. Consider the set $\Sigma = \{\phi_n, : n = 2, \ldots\}$ of UNCAF formulas, where ϕ_n is

$$\forall v_1, \ldots, \forall v_n(\neg(v_1, v_2) \in R \vee \neg(v_2, v_3) \in R\vee, \ldots, \vee\neg(v_n, v_1) \in P).$$

By a well-known theorem (see Richter (1966) and Suzumura (1976)), if a structure (X, R^*, P^*) satisfies these sentences, then it is in T'_{wo}. And if a structure (X, R^*, P^*) is in T'_{wo}, it is clear to see it satisfies these sentences. So $(X, R^*, P^*) \in T'_{wo}$ if and only if it is in T' and satisfies the formulas in Σ. Then, by Theorem 26 $T'_{wo} = T' \cap \mathcal{T}(\Sigma)$ implies that T'_{wo} is falsifiably complete with respect to T', as the formulas in Σ are all UNCAF.

Note that for (X, R^*, P^*) in T', the interpretation of the sentences in Σ is that the strong axiom of revealed preference holds.[17] Note that it is meaningful to talk about a finite data set as "satisfying" a collection of sentences in this case, so long as the sentences do not refer to any constants. This is because there are no function symbols in our language. A data set in this context is a structure for our language ignoring constants. Formally, Afriat's theorem then states that if a finite data set $(D, R^{\mathcal{D}}, P^{\mathcal{D}})$ satisfies the sentences in Σ, there is a structure (X, R^*, P^*) in T_c containing it.

Let (X, R^*, P^*) be a structure in $T'_{wo} \setminus T_c$, and let $\mathcal{D}$ be a finite data set contained in (X, R^*, P^*). It is easy to verify that each of the axioms in Σ are true for $\mathcal{D}$. So, there exists $\mathcal{M} \in T_c$ containing $\mathcal{D}$ by the argument implied by Afriat's theorem.

Since T'_{wo} is falsifiably complete, we conclude that T'_{wo} is the falsifiable closure of T_c with respect to T'. □

7 Other notions of refutability

We are not the first to formalize the notions of falsification and Popper's logical positivism. We discussed the work of Adams, Fagot, and Robinson (1970), Adams (1992) and Pfanzagl, Baumann, and Huber (1971) in the introduction. The excellent book by Luce, Krantz, Suppes, and Tversky (1990) discusses these contributions. Here, we discuss an approach whose formalism is more similar to ours. In a series of papers, Herbert Simon and coauthors (Simon and Groen, 1973; Simon, 1979, 1983, 1985; Rynasiewicz, 1983; Shen and Simon, 1993) discuss a notion of falsifiability, and the formal structure

[17] In first-order logic, the strong axiom is an infinite number of axioms, as we make evident here.

of falsifiable theories. The focus of this work, as we mentioned in the introduction, is on the elimination of theoretical terms.

This literature has based the idea of falsification on the notion of data as a substructure. We now discuss their notion of falsification, and argue that substructures are inadequate as a notion of data. The definition of falsifiability was proposed by Simon and Groen (1973).[18] They intend their definition to capture the theories that can be axiomatized using only universal quantifiers.

A structure $\mathcal{M}$ is *finite* if its domain M is finite.

39 Definition. A theory T is *finitely testable* if there is a structure $\mathcal{M}$ that is not a model of T, and if, for every structure $\mathcal{M}$ that is not a model of T, $\mathcal{M}$ has a finite substructure that is not a model of T.

40 Definition. A theory T is *irrevocably testable* if no model of T has a finite substructure that is not a model of T.

Thus T is finitely and irrevocably testable (FIT) if there is a structure that is not a model of T, and if for every structure $\mathcal{M}$, $\mathcal{M}$ is not a model of T if and only if $\mathcal{M}$ contains a finite substructure that is not a model of T. That is, $\mathcal{M}$ is a model of T if and only if every finite substructure of $\mathcal{M}$ is a model of T. Note that this latter condition also appears in Theorem 30, on relational systems. FIT is the notion of falsifiability used by Simon and Groen. It build on substructures as a notion of data. Note that a relative definition exists: for $T \subseteq T'$, T is FIT with respect to T' if there exists a structure in T' that is not a model of T, and if for every structure $\mathcal{M} \in T'$, $\mathcal{M}$ is not a model of T if and only if $\mathcal{M}$ contains a finite substructure that is not a model of T.

41 Proposition. *If a theory satisfies FIT then it is closed under substructures.*

Proof. Let T satisfy FIT. Let $\mathcal{M}$ be a structure in T. If $\mathcal{M}$ has a substructure that is not in T then this substructure has a finite substructure $\mathcal{B}$ that is not in T. But $\mathcal{B}$ is also a substructure of $\mathcal{M}$, so FIT implies that $\mathcal{M}$ is not in T. It follows that $\mathcal{M}$ cannot have any substructure that is not a model of T. □

By Proposition 41 and the Łoś-Tarski Theorem, FIT implies a universal axiomatization whenever T is a first order theory. The relation between falsifiability and the Łoś-Tarski Theorem is, we hope, clear from our results in Section 5.3.

[18] Rynasiewicz (1983) proposes a different notion, which he calls "finitely strongly falsifiable." One can show that example 42 presents a theory that is falsifiably complete, and closed under substructures, but is not finitely strongly falsifiable.

The following example shows that a theory T may be falsifiably complete with respect to another theory T' [Definition 12], but fail to be FIT (with respect to T'). The example points out that FIT-ness may fail simply because there are no finite substructures of a theory. This can occur for technical reasons related to the definition of substructure.

42 Example. Consider the language $L = \langle 0, q, <, f\rangle$ where q is an unary relation symbol, $<$ is a binary relation symbol, f is a one-place function symbol, and 0 is a constant symbol. Let T' be the class of structures isomorphic to some $\mathcal{M} = (\mathbb{Z}, 0^{\mathcal{M}}, q^{\mathcal{M}}, <^{\mathcal{M}}, f^{\mathcal{M}})$ where $0^{\mathcal{M}}$ is 0 in $\mathbb{Z}$, $<^{\mathcal{M}}$ is a linear order and $x <^{\mathcal{M}} f^{\mathcal{M}}(x)$.

Let T be the class of structures in T' where the formula

$$\forall x \neg q(x)$$

is true. Then by Theorem 26, T is falsifiably complete with respect to T'.

T is also closed under substructures because, if $(\mathbb{Z}, 0^{\mathcal{M}}, q^{\mathcal{M}}, <^{\mathcal{M}}, f^{\mathcal{M}})$ is isomorphic to a model of T and $\mathcal{B}$ is a substructure of $\mathcal{M}$, then $q_{\mathcal{B}}$ coincides with the $q^{\mathcal{M}}$ on $|\mathcal{B}|$.

On the other hand, no model of T' contains any finite substructures. Suppose, to the contrary, that $\mathcal{B}$ is a substructure of $\mathcal{M} \in T'$ and that $|\mathcal{B}|$ is finite. Then $|\mathcal{B}|$ has a largest element $\bar{z}$ according to $<_{\mathcal{B}}$. Note that $f_{\mathcal{B}} = f^{\mathcal{M}}|_{|\mathcal{B}|}$ and $\bar{z} <^{\mathcal{M}} f^{\mathcal{M}}(\bar{z}) = f_{\mathcal{B}}(\bar{z}) \in |\mathcal{B}|$. But $\bar{z}, f_{\mathcal{B}}(\bar{z}) \in |\mathcal{B}|$ and $\bar{z} <^{\mathcal{M}} f_{\mathcal{B}}(\bar{z})$ imply that $\bar{z} <_{\mathcal{B}} f_{\mathcal{B}}(\bar{z})$, which contradicts that $\bar{z}$ was the largest element of $|\mathcal{B}|$.

Consequently, if T were to satisfy FIT with respect to T', it must contain every model of T', which is false. It follows that T does not satisfy FIT with respect to T'.

A theory may satisfy FIT but fail to be falsifiably complete; a simple example involves one unary relation R and theory T axiomatized by $\forall R(x)$.

8 Conclusion

We have developed a theory of the empirical content of an economic theory. The leading examples, throughout the paper, are borrowed from revealed-preference theory; they should be familiar to most economists. We have also shown that the results are applicable to less well-understood theories, and can give new substantive results. In particular, we have illustrated the usefulness of our results by presenting conditions under which theories of multiple-selves in behavioral economics, and theories of preference aggregation in social choice, are falsifiably complete. That is, all its claims are fully testable.

A recurring methodological issue in economics is the argument over unreal assumptions. There is an early literature, sparked by Milton Friedman's 1953 position that the truth of assumptions does not matter. Recent methodological discussions by Rubinstein (2006), Gul and Pesendorfer (2008), Dekel and Lipman (2009), and Gilboa (2009), deal with (among other issues) whether the truth of the "story" behind a theory is relevant. In our results, assumptions and stories do not appear explicitly. They appear implicitly in the specification of concrete theories (see for example the theories in Example 3, and Sections 6.1 and 6.2). This is because we have focused on the testable implications of a theory: an UNCAF axiomatization can be seen as a test for the theory.

However, the framework we have laid out is applicable to the treatment of theoretical objects. We have already mentioned one venue for application using Proposition 29; this result can in fact be applied to study the testable implications of Nash equilibrium or Nash bargaining, something we omitted from the paper because the details are involved and the paper is already long as it is. A second illustration lies in Paul Samuelson's (see Archibald, Simon, and Samuelson (1963)) response to Friedman's position on assumptions. Samuelson effectively counters Friedman by using ideas that we have formalized in our paper. Samuelson makes the point that assumptions matter because either a theory T (described by its "assumptions") is falsifiably complete and thus equivalent to its empirical content, in which case Friedman's point is moot; or it makes non-falsifiable claims, in which case the failure to refute the theory is uninformative about the theory's non-falsifiable claims. In fact, Samuelson argues, by Occam's Razor one should choose the weaker theory, consisting of the empirical content of T (what we have formally termed $\mathrm{fc}(T)$), rather than unnecessary claims in T. Regardless of one's position on the question of realism, we hope that this example shows how our notions may be useful.

Finally, we have studied basic ideas from philosophical positivism. They are seen as naive by some philosophers because researchers may have complicated agendas, and be motivated by their environment, in ways that makes falsification not the focus of their research. Philosophy of science since Popper has therefore focused on the sociology of what drives actual research. We are not expert on these matters, of course, but it seems to us that most economists still find the problem of falsification interesting. In fact, the recent methodological discussions in Gul and Pesendorfer (2008), Dekel and Lipman (2009), and Gilboa (2009), all take for granted that one wants to understand a theory's empirical content (possible exceptions are Hicks (1983) and Rubinstein (2006)). We believe that a formal understanding of empirical content is useful, independently of the complexities involved in the actual production of research.[19]

[19]Gilboa (2009; Chapter 7.3) presents this viewpoint very convincingly.

Appendix A The dual of falsifiable completeness

We have so far discussed falsifiability as a primitive notion, but falsifiability has a dual concept: verifiability. The simplest way to explain these concepts using those we already have is as follows. We can say that a theory T is *verifiably complete* with respect to T' if $T' \backslash T$ is falsifiably complete with respect to T'. Hence, just as falsifiable completeness specifies that all claims of a theory should be falsifiable, verifiable completeness specifies that all claims should be verifiable. Falsifying the complement of a theory is the same as verifying the theory itself–in this sense, falsification and verification are dual.

We can then define the *verifiable interior* of a theory T with respect to T', $\mathrm{vi}_{T'}(T) = T' \backslash \mathrm{fc}_{T'}(T' \backslash T)$. Thus, the verifiable interior of a theory T with respect to T' is the largest subtheory of T which is verifiably complete. It corresponds to the weakest strengthening of the hypotheses for which the theory becomes verifiably complete. Unsurprisingly, the verifiable interior operation is a topological interior, corresponding to the same topology as the falsifiable closure.

Lastly, we can define a sentence to be an ECAF (existential conjunction of atomic formulas) if it is a sentence of the form

$$\exists v_1 \exists v_2 ... \exists v_n (\phi_1 \wedge \phi_2 ... \wedge \phi_n)$$

where each ϕ_i is an atomic formula.

The following result is a trivial consequence of Theorem 26.

43 Theorem. *A theory T is verifiably complete with respect to T' if and only if there exists a set of ECAF sentences, Λ, for which $T = (\bigcup_{\lambda \in \Lambda} \mathcal{T}(\lambda)) \cap T'$.*

We present here a simple example of a theory which is verifiably complete.

44 Example. The example here is one in which we study a private-goods economy, where each individual has her own consumption. We will thus speak of *allocations*. The theory of egalitarian equivalence of some specified allocation, described by Pazner and Schmeidler (1978), asks whether there is some fixed consumption bundle for which each individual is indifferent between her private consumption and the fixed consumption.

To model this, we will suppose that each individual has a preference, and we will consider some fixed allocation; this fixed allocation will be specified in our language by constant symbols.

The language $\mathcal{L}$ involves n binary predicates $R_1, ..., R_n$ and n constant symbols, $c_1, ..., c_n$. The theory that $(c_1, ..., c_n)$ is an egalitarian equivalent allocation is axiomatized by the following sentence:

$$\exists x \bigwedge_{i=1}^{n} (xR_ic_i \wedge c_iR_ix)$$

.

This axiom is immediately seen to be of the ECAF form; hence the theory that $(c_1, ..., c_n)$ is egalitarian equivalent is a verifiably complete theory. This is intuitive, as to verify that the theory holds, one must simply demonstrate the existence of x to which each individual is indifferent.

Appendix B 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 $(x_1, ..., x_m)$.

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

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

1. a set of function symbols $\mathcal{F}$ and positive integers n_f for each $f \in \mathcal{F}$
2. a set of relation symbols $\mathcal{R}$ and positive integers n_R for each $R \in \mathcal{R}$
3. a set of constant symbols $\mathcal{C}$.

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

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

1. a nonempty set M called the *universe* or *domain* of M
2. a function $f^{\mathcal{M}} : M^{n_f} \to M$ for each $f \in \mathcal{F}$

3. a set $R^{\mathcal{M}} \subseteq M^{n_R}$ for each $R \in \mathcal{R}$

4. 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 $f^{\mathcal{M}}$, $R^{\mathcal{M}}$, and $c^{\mathcal{M}}$ are called *interpretations* of the corresponding symbols in the language $\mathcal{L}$.

It is useful to be able to give a meaning to certain relations *across* structures. For example, in our case, we have reason to study both the notion of *substructure* and *isomorphism*. The following makes these precise.

47 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. $\eta(f^{\mathcal{M}}(a_1, ..., a_{n_f})) = f^{\mathcal{N}}(\eta(a_1), ..., \eta(a_{n_f}))$ for all $f \in \mathcal{F}$ and $a_1, ..., a_{n_f} \in M$

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

3. $\eta(c^{\mathcal{M}}) = c^{\mathcal{N}}$ for $c \in \mathcal{C}$.

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

49 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.

50 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.

51 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

52 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 $\mathcal{W}$.

53 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.

54 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).

55 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 $b \in M$, $\mathcal{M} \models \psi(\overline{a}, b)$.

56 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.

57 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

Adams, E. (1992): "On the empirical status of measurement axioms: The case of subjective probability," in *Philosophical and foundational issues in measurement theory*, pp. 53–73. Lawrence Erlbaum Associates.

Adams, E., R. Fagot, and R. Robinson (1970): "On the Empirical Status of Axioms in Theories of Fundamental Measurement," *Journal of Mathematical Psychology*, 7(3), 379–409.

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

Ambrus, A., and K. Rozen (2008): "Revealed conflicting preferences," working paper.

Archibald, G. C., H. A. Simon, and P. A. Samuelson (1963): "Discussion," *The American Economic Review*, 53(2), 227–236.

DeClippel, G., and K. Eliaz (2009): "Reason-Based Choice: A Bargaining Rationale for the Attraction and Compromise Effects," working paper.

Dekel, E., and B. Lipman (2009): "How (not) to do decision theory," working paper.

Dushnik, B., and E. Miller (1941): "Partially ordered sets," *American Journal of Mathematics*, 63(3), 600–610.

Friedman, M. (1953): "The Methodology of Positive Economics," in *Essays in positive economics*. The University of Chicago Press, Chicago Illinois.

Fudenberg, D., and D. Levine (2006): "A dual-self model of impulse control," *American Economic Review*, 96(5), 1449–1476.

Gilboa, I. (2009): *Theory of decision under uncertainty*. Cambridge University Press.

Green, J. R., and D. Hojman (2008): "Choice, rationality, and welfare measurement," working paper.

Gul, F., and W. Pesendorfer (2008): "The case for mindless economics," *The foundations of Positive and normative Economics: A handbook*, pp. 3–42.

Hicks, J. (1983): "A discipline not a science," in *Classics and Moderns: Collected Essays on Economic Theory Volume 3*, pp. 365–375.

Kalai, G. (2004): "Social indeterminacy," *Econometrica*, 72, 1565–1581.

Kelley, J. L. (1955): *General Topology*. D. Van Nostrand Company, Inc.

Levy, A. (2002): *Basic set theory*. Dover.

Luce, R., D. Krantz, P. Suppes, and A. Tversky (1990): *Foundations of measurement, Vol. 3: Representation, axiomatization, and invariance*. Academic.

Manzini, P., and M. Mariotti (2007): "Sequentially rationalizable choice," *American Economic Review*, 97(5), 1824–1839.

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

McGarvey, D. C. (1953): "A theorem on the construction of voting paradoxes," *Econometrica*, 21(4), 608–610.

O'Donoghue, T., and M. Rabin (1999): "Doing It Now or Later," *The American Economic Review*, 89(1), 103–124.

Pazner, E., and D. Schmeidler (1978): "Egalitarian equivalent allocations: A new concept of economic equity," *The Quarterly Journal of Economics*, 92(4), 671–687.

Pfanzagl, J., V. Baumann, and H. Huber (1971): *Theory of Measurement*. Physica-Verl.

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

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

Rubinstein, A. (2006): "Dilemmas of an economic theorist," *Econometrica*, 74(4), 865–883.

Rynasiewicz, R. A. (1983): "Falsifiability and the Semantic Eliminability of Theoretical Languages," *The British Journal for the Philosophy of Science*, 34(3), 225–241.

Sen, A. (1969): "Quasi-transitivity, rational choice and collective decisions," *The Review of Economic Studies*, 36(3), 381–393.

Shen, W.-M., and H. A. Simon (1993): "Fitness Requirements for Scientific Theories Containing Recursive Theoretical Terms," *The British Journal for the Philosophy of Science*, 44(4), 641–652.

Simon, H. A. (1979): "Fit, Finite, and Universal Axiomatization of Theories," *Philosophy of Science*, 46(2), 295–301.

——— (1983): "Fitness Requirements for Scientific Theories," *The British Journal for the Philosophy of Science*, 34(4), 355–365.

——— (1985): "Quantification of Theoretical Terms and the Falsifiability of Theories," *The British Journal for the Philosophy of Science*, 36(3), 291–298.

Simon, H. A., and G. J. Groen (1973): "Ramsey Eliminability and the Testability of Scientific Theories," *The British Journal for the Philosophy of Science*, 24(4), 367–380.

Sprumont, Y. (2001): "Paretian quasi-orders: the regular two-agent case," *Journal of Economic Theory*, 101(2), 437–456.

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.

Trotter, W. T. (1992): *Combinatorics and Partially Ordered Sets*. The Johns Hopkins University Press.

Varian, H. R. (1982): "The Nonparametric Approach to Demand Analysis," *Econometrica*, 50(4), 945–974.

Vaught, R. (1954): "Remarks on universal classes of relational systems," *Indagationes Mathematicae*, 16, 589–591.