To Infinity and Beyond: Scaling Economic Theories via Logical Compactness*

Yannai A. Gonczarowski† Scott Duke Kominers‡ Ran I. Shorrer§

August 15, 2020

Abstract

Many economic-theoretic models incorporate finiteness assumptions that, while introduced for simplicity, play a real role in the analysis. Such assumptions introduce a conceptual problem, as results that rely on finiteness are often implicitly nonrobust; for example, they may depend upon edge effects or artificial boundary conditions. Here, we present a unified method that enables us to remove finiteness assumptions, such as those on market sizes, time horizons, and datasets. We then apply our approach to a variety of matching, exchange economy, and revealed preference settings.

The key to our approach is *Logical Compactness*, a core result from Propositional Logic. Building on Logical Compactness, in a matching setting, we reprove large-market existence results implied by Fleiner's analysis, and (newly) prove both the strategy-proofness of the man-optimal stable mechanism in infinite markets and an infinite-market version of Nguyen and Vohra's existence result for near-feasible stable matchings with couples. In a trading-network setting, we prove that the Hatfield et al. result on existence of Walrasian equilibria extends to infinite markets. In a dynamic matching setting, we prove that Pereyra's existence result for dynamic two-sided matching markets extends to a doubly infinite time horizon. Finally, beyond existence and characterization of solutions, in a revealed-preference setting we reprove Reny's infinite-data version of Afriat's theorem and (newly) prove an infinite-data version of McFadden and Richter's characterization of rationalizable stochastic datasets.

*A one-page abstract of this paper appeared in the Proceedings of the 21st ACM Conference on Economics and Computation (EC 2020). The authors thank David Ahn, Bob Anderson, Morgane Austern, Chris Chambers, Yunseo Choi, Henry Cohn, Piotr Dworczak, Drew Fudenberg, Wayne Gao, Jerry Green, Joseph Halpern, Ron Holzman, Ravi Jagadeesan, M. Ali Khan, David Laibson, Rida Laraki, Bar Light, Ce Liu, George Mailath, Michael Mandler, Paul Milgrom, Ankur Moitra, Yoram Moses, Marek Pycia, John Rehbeck, Phil Reny, Joseph Root, Dov Samet, Chris Shannon, Sergiy Verstyuk, Shing-Tung Yau, and Bill Zame for helpful comments. Gonczarowski was supported in part by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities; his work was supported in part by ISF grants 1435/14, 317/17, and 1841/14 administered by the Israeli Academy of Sciences; by the United States–Israel Binational Science Foundation (BSF grant 2014389); and by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 740282), and under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122. Additionally, Kominers gratefully acknowledges the support of the National Science Foundation (grant SES-1459912), as well as the Ng Fund and the Mathematics in Economics Research Fund of the Harvard Center of Mathematical Sciences and Applications. Shorrer was supported by a grant from the United States–Israel Binational Science Foundation (BSF grant 2016015).

†Microsoft Research — *E-mail*: yannai@gonch.name. Parts of the work of Gonczarowski were carried out while at the Hebrew University of Jerusalem and at Tel Aviv University.

‡Entrepreneurial Management Unit, Harvard Business School; Department of Economics, Center of Mathematical Sciences and Applications, and Center for Research on Computation and Society, Harvard University; and National Bureau of Economic Research — *E-mail*: kominers@fas.harvard.edu.

§Department of Economics, Penn State University — *E-mail*: shorrer@psu.edu.

1 Introduction

In microeconomic theory, we frequently make finiteness assumptions for simplicity and/or tractability—and those assumptions can play a real role in our analysis. Of course, the real world is itself finite, so there is in some sense no "loss" from assuming finiteness in our models. But there is sometimes a conceptual problem: if our understanding of economic theory hinges on finiteness, then our models may not quite tell the whole story. For example:

- If a game-theoretic finding is true only when the set of agents is finite, there is an implicit discontinuity, possibly relying on an edge effect or a specific starting condition that may not be robust to small frictions or perturbations (e.g., the existence of a highest-surplus match).[1] Thus finite-market results that extend to infinite markets are in some sense more foundational.

- Dynamic games serve to model long-run behavior and steady-states—representing interactions that will be repeated over and over with no fixed start or end time. We thus turn to infinite time horizons instead of long finite horizons in order to avoid behavioral artefacts driven by time "starting" or "ending" at a fixed time.

- And in decision theory, revealed preference analysis seeks to understand what we can infer about an agent from his or her choice behavior. While a list of observed choices is always finite, we like to reason about how an agent's choice would behave on arbitrary input data—and this by nature requires conjecturing about behavior over an infinite dataset. Furthermore, theorizing about observing infinite datasets lets us characterize limitations in inference about agents' preferences that are inherent—as opposed to just imposed by finiteness of data.

Yet many of our results that use finiteness seem like they should logically scale to infinite settings, as well.[2] In this paper, we show that the preceding intuition is precisely—and in fact, verbatim—correct, at least for a number of canonical results in matching, trading networks, and even decision theory.[3] In each case, we show how to carry over results from finite models to infinite ones—in each of the senses of "infinite" just described—by way of *Logical Compactness*, a central result in the theory of Propositional Logic, which (roughly) states that an infinite set of individually finite logical statements can be made consistent if and only if every finite subset can.[4]

In some of the settings we consider, certain results have already been lifted from finite models to infinite ones. But heretofore such liftings have required adapting—and in some cases completely restructuring—the core of each finite-case proof using specialized, setting-specific tools. Extending results in matching theory to infinite markets, for example, has often relied on versions of Tarski's theorem. And in game theory, scaling Nash's equilibrium theorem typically relies on specialized generalizations of the Brouwer fixed point theorem, such as the Schauder fixed-point theorem. Reny's (2015) generalization of Afriat's theorem to infinite datasets required a completely novel construction. Yet despite the range of settings we consider, our approach is more or less uniform throughout: We show how to rephrase each economic theory setting under consideration in terms of a collection of individually finite logical statements; Logical Compactness then enables us to directly translate results from finite instantiations of the model to infinite ones. In this approach,

[1] For an example of a different kind of discontinuity—between a finite and a continuum setting—see the work of Miralles and Pycia (2015), showing that the continuum economy sometimes rules out important phenomena that are observed in finite markets that converge to it.

[2] As a side note: We sometimes also turn to infinite models when we cannot solve their finite analogues—for example, to smooth out integer effects. But that is not our focus here.

[3] In Appendix C, we show how our methods also apply to games on graphs.

[4] We give a formal statement of the Compactness Theorem for Propositional Logic in Section 2.

Logical Compactness functions in a "black-box" manner: we lift theorems with similar *statements* using similar techniques, even if the known proofs of those theorems are in fact very different.

For example, we encode a (one-to-one) stable matching in terms of a set of Boolean variables $\{\texttt{matched}_{(m,w)}\}$ that are TRUE when man m is matched to woman w. (We further elaborate on this example in Section 3.) To make sure that what we encode is a feasible matching, we start by adding, for each agent i and pair of agents j, j' from the opposite side, a logical statement requiring that i is matched to at most one of j or j'. To make sure that the matching is individually rational, if some man m finds woman w unacceptable, or vice versa, then we introduce a logical statement that requires $\texttt{matched}_{(m,w)}$ to be FALSE. Finally, to ensure the stability of the encoded matching, we introduce logical statements ruling out blocking pairs, requiring that if $\texttt{matched}_{(m,w)}$ is FALSE even though m and w find each other acceptable, then we must either have $\texttt{matched}_{(m,w')}$ TRUE for some w' that m prefers to w or $\texttt{matched}_{(m',w)}$ TRUE for some m' that w prefers to m. A logically consistent assignment of truth values to the $\texttt{matched}_{(m,w)}$ variables then corresponds exactly to a stable matching, and vice versa.

The logical formulation just described captures finite and infinite matching models equally: the only difference that arises is in the cardinalities of the set of variables and of the set of logical statements. When the set of agents is infinite, the associated set of logical statements is infinite as well. But so long as agents' preferences reflect (possibly infinite) rank-order lists, each of the logical statements we use is individually finite—so that every finite subset of the infinite set of logical statements corresponds to an ordinary, finite matching problem. Known existence results for finite matching models (Gale and Shapley, 1962) thus give us a consistent logical solution for every finite subset. Logical Compactness then yields a consistent solution—and hence existence of stable outcomes—in the infinite model (even with infinite preference lists).

While existence of stable one-to-one matchings in infinite economies has been known since the work of Fleiner (2003), our approach to proving existence extends to settings outside the scope of Fleiner's result and methods. Indeed, we extend our model to show how to lift a recent result of Nguyen and Vohra (2018) on matching with couples to infinite markets. Specifically, we use Logical Compactness to prove that in an infinite doctor–hospital matching market with couples, it is possible to find a small perturbation of hospitals' capacities such that at least one stable matching exists—so long as couples' preferences satisfy a mild "downward closure" condition. This result directly extends the analogous finite-market result of Nguyen and Vohra (2018); consequently, the size of the capacity perturbation needed exactly corresponds to that Nguyen and Vohra (2018) found (in particular, the perturbation is finite). Here, the fact that Logical Compactness is agnostic to the proof of the finite-market result is particularly powerful, as the result of Nguyen and Vohra (2018) relies on Scarf's Lemma and other tools far removed from typical existence proofs in matching theory, such as deferred acceptance and the methods Fleiner (2003) used.

Our matching model formulation also lets us generalize man-optimality and strategy-proofness results for stable matching to infinite markets—although the arguments required to apply Logical Compactness in those cases are more subtle.[5] Our strategy-proofness result—which is novel to the present work—is perhaps particularly surprising because standard proofs of strategy-proofness rely on versions of the Lone Wolf/Rural Hospitals Theorem (Roth, 1984), which Jagadeesan (2018b) has shown *fails* in our setting. Our method thus shows that strategy-proofness of the man-optimal stable matching mechanism is in some sense more robust/fundamental than the Lone Wolf/Rural Hospitals result that is typically used to prove it.[6] Moreover, in the infinite markets we obtain under

[5]Like with existence, the man-optimality result was originally proven by Fleiner (2003) using fixed-point methods.

[6]There are a number of matching settings in which the Rural Hospitals Theorem fails but strategy-proofness is still obtained (see, e.g., Kamada and Kojima, 2015; Hatfield and Kominers, 2019); however, to our knowledge, in all of these settings the proof of strategy-proofness still relies on a version of the Rural Hospitals Theorem in an auxiliary market.

Logical Compactness, agents in a sense "maintain their mass," which is particularly economically appealing because strategic issues remain in full force.[7]

We also use Logical Compactness to obtain existence results for Walrasian equilibria in large trading networks. In that case (and unlike in matching), continuous action/price spaces mean that the equilibrium object is not itself locally finite—even in finite markets. Thus, we use Logical Compactness to obtain existence of approximate equilibria in the infinite market instead; we then translate those approximate equilibria into full equilibria by way of a diagonalization argument.

Next, we show how to use Logical Compactness to extend finite horizons to infinite ones: We consider the dynamic matching setting of Pereyra (2013), in which teachers arrive and depart at different periods, and must be matched stably subject to a "tenure" rule that gives teachers the right to remain in the schools they are assigned to before the next set of teachers arrives. Pereyra (2013) relies on a "period 0" to fix the initial allocation to some stable matching, and then progresses to future periods recursively; using Logical Compactness, we can dispense with the assumption of a "period 0" to obtain existence results in an infinite-past-horizon model, which is perhaps more appropriate as a representation of a steady state (see, e.g., Öry, 2016; Clark, Fudenberg, and Wolitzky, 2019).[8]

We illustrate moreover that Logical Compactness has applications outside of the existence and characterization of solutions. In particular, we show a few ways our approach can be used in decision-theoretic settings: Szpilrajn's Extension Theorem—whose variants are central to classical results in consumer theory—follows immediately from Logical Compactness and an easy-to-prove finite case. In the case of Afriat's theorem, an approximate version follows immediately from the finite case by way of Logical Compactness; we impose on these approximations certain restrictions consistent with the guarantee of Afriat's theorem, and this enables us to exactly recover the infinite-data generalization first proven by Reny (2015). Then, we use Logical Compactness to develop a novel generalization of the McFadden and Richter (1971, 1990) characterization of rationalizable stochastic datasets.

In all of the settings we consider, the key benefit of our approach is that once we have a logical formulation, Compactness arguments permit us to reason solely about finite models in their regular language of analysis. Logical Compactness allows us to obtain new results on problems of existing interest such as matching with couples, dynamic matching, and stochastic choice; at the same time, Compactness helps make certain key conditions in infinite models more transparent (e.g., quasiconcavity in the result of Reny (2015); see Section 6.1). And Logical Compactness is applicable in cases (like our strategy-proofness result) where it is not at all clear *ex ante* whether a limit-based formulation is even possible.

1.1 Relation to Other Mathematical Approaches

It is important to note that logical statements can be translated into closed sets in an application-specific topological (product) space, in which setting Logical Compactness follows from Tychonoff's theorem on topological compactness. In other words, whatever can be proved using Logical Compactness can also be proved via topology and Tychonoff's theorem. However, to non-experts, the resulting proof would be harder to directly formulate and harder to verify. And while topological compactness or the language of nets are stronger and more general approaches, in the domains we study, they often introduce technical issues that can render arguments incorrect in subtle ways (e.g.,

[7] In continuum-limit models of large markets, by contrast, strategic issues often disappear because agents become measure-0 (see, e.g., Azevedo and Budish, 2018).

[8] Clark, Fudenberg, and Wolitzky (2019) refer to time in such a model as "doubly infinite," and indeed note specifically such a model is conceptually useful in excluding strategies that condition on calendar time.

matchings may converge to an object that is not a matching). We therefore view the methodological part of our contribution as introducing a unifying approach that is simple and intuitive to work with, and that does not require us to look for the "right" topological space or apply topological reasoning directly.[9]

Working with Logical Compactness reduces reasoning about infinite problems to reasoning about finite problems—and we can think about those finite problems directly in a natural language. It is of course nontrivial to construct a logical formulation of a given theory model, and to make sure that all the appropriate local finiteness conditions hold up. But nevertheless, our approach moves the challenge of scaling economic theories away from technical details such as topological structure and measure-theoretic concerns.

We have also been asked about alternative proofs using first-order logic and the transfer principle. Propositional logic is a special case of first-order logic. Importantly, it does not use quantifiers (i.e., $\forall$ and $\exists$). Working with fully fledged first-order logic (i.e., where the domain of the model corresponds, e.g., to the set of agents in the market, or more generally simply with a non-empty domain and with quantifiers over that domain) would not have allowed us to fix the set of objects (e.g., men and women) in our economic model. In fact, by the (upward) Löwenheim–Skolem theorem, if a first-order theory has an infinite model (a model with an infinite domain) then it has a model of any larger cardinality, which implies that first-order theories cannot bound the cardinality of their infinite models. But in many of our applications, we do want to work with a specific infinite set (in this case, the set of players) and not "allow" the tools used in the proof to create new elements (in this case, additional players) to satisfy our requirements. For this reason we work with purely propositional theories (or equivalently, quantifier-free theories whose models have empty domains and where the primitives are constants—nullary relations that already explicitly encode the elements of our economic model). By contrast, the transfer principle is used, roughly speaking, to prove statements in nonstandard models with augmented domains—sometimes also using first-order compactness in those models—and then transfer these results to the standard models of interest, which have more restricted (yet still rich) domains.

Finally, we emphasize that our approach is not without limitations (see Section 3.4 for a discussion). However, the limitations of methods based on Logical Compactness in some cases give insight into the boundaries of when finiteness assumptions can be relaxed. For example, as we discuss in Section 3.4, if we try to construct a Compactness formulation of the Lone Wolf/Rural Hospitals Theorem for matching markets—which Jagadeesan (2018b) has shown fails in infinite settings—we quickly run into a need for statements outside of the boundaries of first-order logic, such as the ability to condition one model on another. Likewise, for certain types of matching-theoretic efficiency results, we find ourselves needing individually infinite formulae to rule out infinite efficiency-improving trading cycles among agents, which again takes us outside the scope of our methods—and indeed, Choi (2020) has subsequently shown that precisely because of the presence of infinite cycles, some classical efficiency results do not extend to infinite markets.[10]

[9] Once a proof is derived using Logical Compactness, it is of course possible to then translate it to a topological statement and attempt to achieve greater generality, if/when such generality is of interest.

[10] These features manifest not only when considering infinite markets, but also in our other use cases, of infinite time and infinite datasets. In our generalization of the dynamic matching result of Pereyra (2013) to doubly infinite time horizons in Section 5, making the logical formulae in our proof individually finite necessitated the introduction of a finite presence condition, which helped us flush out a counterexample showing that this condition cannot be dispensed with. Meanwhile, in our reproof of the result of Reny (2015) in Section 6, it may seem that the ability to only construct quasiconcave rather than concave utility functions stems from a discretization argument necessary for us to make the logical formulae in our proof individually finite; nonetheless, it is well known that concavity cannot be guaranteed in this context, and our Logical Compactness argument reproves the result under the same conditions in which Reny (2015) proved it.

1.2 Related Literature

Infinite models are used frequently in game theory as a way of representing limit—or "large"—markets. Many of these models work with either discrete infinite markets (Fleiner, 2003; Kircher, 2009; Jagadeesan, 2018a) or a limit of finite markets (Immorlica and Mahdian, 2005; Kojima and Pathak, 2009; Ashlagi et al., 2014; Miralles and Pycia, 2015); we work with similar limit settings, but use Compactness to sidestep many of the challenges in analyzing those models.[11] Our existence and structural results for large matching markets (but not our strategy-proofness result) are implicitly covered by the main result of Fleiner (2003), who introduced a fixed-point characterization of stable matchings that also holds in infinite markets. Our strategy-proofness result, meanwhile, generalizes the more restricted result of Jagadeesan (2018b), who proved strategy-proofness of the man-optimal stable matching mechanism in markets with countably many agents under a "local finiteness" condition (which we avoid). Subsequent to our work, Choi (2020) used the methods we introduce here to lift several other classic results from finite matching markets to infinite ones: group strategy-proofness (for finite coalitions); a comparative static on the entry of new agents; and the respect for improvements theorem.

In Section 6, we generalize a finite-data theorem for stochastic choice originally due to McFadden and Richter (1971, 1990) to cover infinite datasets.[12] The importance of infinite data in the McFadden and Richter (1971, 1990) setting was highlighted by Cohen (1980) and McFadden (2005). McFadden (2005), in particular, showed how to extend the McFadden and Richter (1971, 1990) result to an infinite setting different than ours by either weakening the concept of rationalizability or by imposing topological structure—neither of which we require in our generalization.

Our proofs of any of the known results that we reprove—and furthermore our proofs of all of the novel results of this paper—are based on one common, principled approach. Thus our exercise here is in some sense similar to that of Blume and Zame (1994), who unified our understanding of perfect and sequential equilibria by way of the Tarski–Seidenberg Theorem. Our work is analogously similar to papers that have used nonstandard analysis to refine and scale results in economic theory (see, e.g., Anderson, 1978; Brown and Khan, 1980; Anderson, 1991; Khan, 1993; Halpern, 2009, 2010; Halpern and Moses, 2016). Our result for converting a dynamic game with a finite start time into an "ongoing" dynamic game with neither start nor end also connects our work to the broad literature on infinite-horizon games (see, e.g., Fudenberg and Levine, 2009).

Methodologically, to our knowledge, we are the first to use the Compactness Theorem for Propositional Logic as a general tool for reinterpreting and scaling results in economics. It is worth mentioning within this context, though, the work of Holzman (1984), who used Compactness to relax topological conditions in Fishburn (1984). Chambers et al. (2014) used compactness in first-order (rather than propositional) logic to formalize the notion of the empirical content of a model; while Chambers et al. (2014), like some of our results us, study applications to revealed preferences theory (see also Chambers et al., 2017), they deal with different questions from us, and use different techniques. Halpern (2010) also uses first-order (rather than propositional) compactness, however within the scope of nonstandard analysis, to prove results about nonstandard probabilities in preferences. Feinberg (2000) gives a syntactic characterization of common priors when the state space is finite or compact under a suitable topology, and uses the equivalence of logical and

[11] A second class of large-market models features continua of agents, with each agent having a negligible contribution to the overall market (see, e.g., Aumann and Shapley, 1974; Gretsky et al., 1992, 1999; Kaneko and Wooders, 1986; Azevedo et al., 2013; Nöldeke and Samuelson, 2018; Azevedo and Budish, 2018; Greinecker and Kah, 2019). A series of recent papers has introduced models that mix between the two styles by featuring countably many "large" agents that can each match with a continuum of "small" agents (see., e.g., Azevedo and Leshno, 2016; Azevedo and Hatfield, 2015; Jagadeesan, 2018a; Che et al., 2019; Fuentes and Tohmé, 2019).

[12] For a recent example of a treatment of infinite datasets see Aguiar et al. (2020).

topological compactness to phrase this compactness assumption on the state space syntactically as well. Finally, Bjorndahl et al. (2014) use compactness in a very different way: they define a general class of games in which an agent's utility depends on which formulae in a certain language are true, and show that if the language is compact then rationalizable strategies are guaranteed to exist. More generally, we follow in a powerful tradition of using logic to organize and extend ideas in game theory, started by the work of Blume and Zame (1994), and continued by Arieli and Aumann (2015) and Hellman and Levy (2019).[13]

Meanwhile, Logical Compactness is frequently used to extend existence results in mathematics from finite settings to infinite ones. For example, de Bruijn and Erdős (1951) and Halmos and Vaughan (1950) respectively used Compactness to derive infinite-graph versions of graph coloring results and Hall's marriage theorem. However, those applications rely on local finiteness conditions—specifically, finite degree—that we are able to avoid here (at least in our matching settings). Moreover, unlike in standard graph-theoretic applications, we manage to use Logical Compactness arguments to prove results about uncountable and continuous objects (such as utility functions and Walrasian prices), as well as complex characterization results (e.g., strategy-proofness) that go beyond existence results and are not inherently topological.

Finally, we note that in addition to the fact that Hellman and Levy (2019) use (different) tools from mathematical logic, their paper is also somewhat conceptually related to ours: while our paper lifts certain "finite results" to "infinite results," their paper lifts certain "countably infinite results" to "*un*countably infinite ones." Specifically, they give sufficient conditions to lift certain existence results that are known to hold whenever there are countably many possible *states of the world* into scenarios with uncountably many possible states of the world. Their results are incomparable to any of our results, and even to our existence-in-large-market results, first because they always assume that the number of agents is finite (an infinite number of agents, even with only two possible types for each, would already result in an uncountably infinite set of possible states of the world to begin with), and second, because they require that the theorems that they lift be already known to hold for the countably infinite, rather than only the finite, case.

1.3 Outline of the Paper

The remainder of this paper is organized as follows. Section 2 introduces preliminaries from Propositional Logic, states the Compactness Theorem, and demonstrates how Compactness works by deriving an immediate proof of Szpilrajn's Extension Theorem (Szpilrajn, 1930). Section 3 serves as a "warm up" by giving a fairly simple illustration of our approach—using Logical Compactness to (re-)prove the existence of stable outcomes in infinite, one-to-one matching markets; this section also discusses limitations of our approach, as well as gives a few words of caution regarding its usage. Section 4 demonstrates the potential of Logical Compactness to remove finiteness assumptions on market sizes, by proving structural and strategy-proofness results for infinite matching markets; by generalizing the result of Nguyen and Vohra (2018) on near-feasible stable matchings with couples to infinite markets; and by proving the existence of Walrasian equilibria in infinite trading networks. Section 5 demonstrates the potential of Logical Compactness to remove finiteness assumptions on time, generalizing the dynamic matching framework of Pereyra (2013) to a doubly infinite time horizon. Section 6 demonstrates the potential of Logical Compactness to remove finiteness assumptions on data, by using Logical Compactness to reprove Reny's generalization of Afriat's theorem and to give a novel generalization of the McFadden and Richter (1971, 1990)

[13]See also the concise proof of Zermelo's theorem described by Maschler et al. (2013) and attributed to Abraham Neyman.

characterization for rationalizability of stochastic choice datasets. Section 7 concludes. Omitted proofs and further applications are presented in the appendix.

2 Propositional Logic Preliminaries

In this section we introduce the Compactness Theorem for Propositional Logic, after quickly reviewing the necessary definitions required to state it.[14] Because the machinery of Propositional Logic can be unfamiliar, we give running examples throughout. Then, to illustrate how we apply Propositional Logic concepts, we use the Compactness Theorem to give a concise proof of Szpilrajn's Extension Theorem, a result on orderings that is used throughout decision theory.

In Propositional Logic, we work with a set of Boolean variables, and study the truth values of statements—called formulae—made up of those variables. We construct formulae by conjoining variables with simple logical operators such as OR, NOT, and IMPLIES. Variables are abstract, and do not have meaning on their own—but we can imbue them with "semantic" meaning by introducing formulae that reflect the structure of economic (or other) problems. Once given semantic meaning, the truth or falsity of statements in our Propositional Logic model imply the corresponding results in the associated economic model.

We start by formalizing the idea of *(well-formed propositional) formulae.* To define the set of formulae at our disposal, we first define a basic (finite or infinite) set of *atomic formulae*, which serve the role of (Boolean) *variables.* Atomic formulae are in some sense the primitives (or basic units) of a Propositional Logical formulation; in each section of this paper we will have a different set of atomic formulae built around the economic setting we are modeling.

Once we have defined a (finite or infinite) set V of atomic formulae, we can define the set of all well-formed formulae inductively:

- Every atomic formula $\phi \in V$ is a well-formed formula.
- '$\neg\phi$' is a well-formed formula for every well-formed formula ϕ.
- '$(\phi\vee\psi)$', '$(\phi\wedge\psi)$', '$(\phi \to \psi)$', and '$(\phi \leftrightarrow \psi)$' are well-formed formulae for every two well-formed formulae ϕ and ψ.

Example. *We could start, for example, with a set of four atomic formulae $V = \{\mathtt{P}, \mathtt{Q}, \mathtt{R}, \mathtt{S}\}$. With the atomic formulae V, each of the following is a well-formed formula:*

$$\text{`}\mathtt{P}\text{'}, \tag{1}$$
$$\text{`}(\mathtt{P} \vee \mathtt{Q})\text{'}, \tag{2}$$
$$\text{`}\neg(\mathtt{P} \wedge \mathtt{Q})\text{'}, \tag{3}$$
$$\text{`}((\mathtt{P} \wedge \mathtt{R}) \to \mathtt{S})\text{'}. \tag{4}$$

We sometimes abuse notation by omitting parentheses and writing, e.g., '$\phi \vee \psi \vee \xi$' when any arbitrary placement of parentheses in the formula (e.g., '$((\phi \vee \psi) \vee \xi)$' or '$(\phi \vee (\psi \vee \xi))$') will make do for our analysis. We will sometimes even abuse notation by writing, e.g., '$\bigvee_{i=1}^{10} \phi_i$' to mean '$\phi_1 \vee \phi_2 \vee \cdots \vee \phi_{10}$' (we will once again do so only when the precise placement of omitted parentheses is of no consequence to our analysis).

We note that while well-formed formulae can be arbitrarily long, each well-formed formula is always finite in length. Thus, for example, a disjunction '$\phi_1 \vee \phi_2 \vee \cdots$' of infinitely many formulae

[14] For a more in-depth look at Propositional Logic primitives and at the Compactness Theorem, see a textbook on Mathematical Logic (e.g., Marker, 2006).

is *not* a well-formed formula. We will therefore take special care when we claim that a formula of the form, e.g., '$\bigvee_{\phi \in \Psi} \phi$' is a well-formed formula, as this is true only if Ψ is a finite set of formulae.

Example. *In particular, for a countably infinite set of atomic formulae $V = \{\mathtt{P}_n\}_{n=1}^{\infty}$ and $X \subseteq \mathbb{N}$,*

$$`\bigvee_{n \in X} \mathtt{P}_n'$$

is a well-formed formula if and only if X is finite.

A *model* is a mapping from the set V of all variables (atomic formulae) to Boolean values, so each variable is mapped either to being TRUE or to being FALSE. A model also induces a *truth value* for every nonatomic formula, defined inductively as follows:

- '$\neg\phi$' is TRUE iff ϕ is FALSE;
- '$(\phi \vee \psi)$' is TRUE iff either or both of ϕ and ψ is TRUE;
- '$(\phi \wedge \psi)$' is TRUE iff both ϕ and ψ are TRUE;
- '$(\phi \to \psi)$' is TRUE iff either ϕ is FALSE or ψ is TRUE or both (that is, '$(\phi \to \psi)$' is FALSE only if both ϕ is TRUE and ψ is FALSE); and
- '$(\phi \leftrightarrow \psi)$' is TRUE iff ϕ and ψ are either both TRUE or both FALSE.

Example. *Given the concept of truth values, we can reinterpret the formulae* (1)–(4) *as follows:*

$$`\mathtt{P}' \qquad \text{“}\mathtt{P}\ \textit{[IS TRUE]}\text{”}, \tag{1}$$
$$`(\mathtt{P} \vee \mathtt{Q})' \qquad \text{“}\mathtt{P}\ \textit{OR}\ \mathtt{Q}\ \textit{[IS TRUE]}\text{”}, \tag{2}$$
$$`\neg(\mathtt{P} \wedge \mathtt{Q})' \qquad \text{“}\textit{NOT (}\mathtt{P}\ \textit{AND}\ \mathtt{Q}\ \textit{[ARE BOTH TRUE])}\text{”}, \tag{3}$$
$$`((\mathtt{P} \wedge \mathtt{R}) \to \mathtt{S})' \qquad \text{“}\mathtt{P}\ \textit{AND}\ \mathtt{R}\ \textit{[BOTH BEING TRUE], IMPLIES}\ \mathtt{S}\ \textit{[BEING TRUE]}\text{”}. \tag{4}$$

The formula in (2) *is TRUE in a model if and only if either* '$\mathtt{P}$' *or* '$\mathtt{Q}$' *(or both) are TRUE in that model; the formula in* (3) *is TRUE in a model unless both* '$\mathtt{P}$' *and* '$\mathtt{Q}$' *are TRUE in that model; and the formula in* (4) *above is TRUE in that model unless both* '$\mathtt{P}$' *and* '$\mathtt{R}$' *are TRUE in that model while* '$\mathtt{S}$' *is FALSE in that model.*

We say that a formula is *satisfied* by a model if it is TRUE under that model. We say that a (possibly infinite) set of formulae is *satisfied* by a model if every formula in the set is satisfied by the model. We say that a (possibly infinite) set of formulae is *satisfiable* if it is satisfied by some model.

Clearly, if a (finite or infinite) set of formulae Φ is satisfiable, then every subset of Φ is also satisfiable (by the same model), and in particular every *finite* subset of Φ is satisfiable; the *Compactness Theorem for Propositional Logic* gives a surprising and nontrivial converse to this statement.

Theorem 2.1 (The Compactness Theorem for Propositional Logic)**.** *A set of formulae Φ is satisfiable if (and only if) every* ***finite*** *subset $\Phi' \subseteq \Phi$ is satisfiable.*

2.1 Illustration: Szpilrajn's Extension Theorem

Now, to illustrate how we apply the concepts just introduced, we use the Compactness Theorem to give a concise proof of the following result known as Szpilrajn's Extension Theorem.

Theorem 2.2 (Szpilrajn's Extension Theorem)**.** *Let X be a set (of any cardinality). Every strict partial order on X can be extended to a total order.*

Variants of Theorem 2.2 are used to prove many key results in decision theory, such as the sufficiency of the strong axiom of revealed preferences for the existence of rationalizing preferences. Such variants are customarily proven using Zorn's Lemma (e.g., Richter, 1966; Duggan, 1999; Mas-Colell, Whinston, and Green, 1995, Proposition 3.J.1; Chambers and Echenique, 2016, Theorems 1.4 and 1.5). While Logical Compactness, like Zorn's Lemma, relies on some variant of the Axiom of Choice, we suspect that the proof we present here may complement the standard approach. In particular, our argument may in some ways be more accessible to students than the traditional proof because it avoids the "overhead" of understanding the full statement of Zorn's Lemma.[15]

Proof of Theorem 2.2. We construct a set of variables V by defining a variable ${}_a\mathtt{gt}_b$ for each pair of distinct $a, b \in X$. Given a strict partial order $>_X$ over X, we can define a set $\Phi_{>_X}$ of formulae consisting of:

- for every distinct $a, b \in X$ such that $a >_X b$, the formula '${}_a\mathtt{gt}_b$';
- for every distinct $a, b \in X$, the formula '${}_a\mathtt{gt}_b \vee {}_b\mathtt{gt}_a$';
- for every distinct $a, b \in X$, the formula '$\neg({}_a\mathtt{gt}_b \wedge {}_b\mathtt{gt}_a)$';
- for every distinct $a, b, c \in X$, the formula '$({}_a\mathtt{gt}_b \wedge {}_b\mathtt{gt}_c) \rightarrow {}_a\mathtt{gt}_c$'.

(Astute readers will notice that the preceding formulae correspond exactly with (1)–(4) upon taking $\mathtt{P} = {}_a\mathtt{gt}_b$, $\mathtt{Q} = {}_b\mathtt{gt}_a$, $\mathtt{R} = {}_b\mathtt{gt}_c$, and $\mathtt{S} = {}_a\mathtt{gt}_c$.)

With the structure just described, the variable ${}_a\mathtt{gt}_b$ has the interpretation "a is greater than b (under the order $>_X$)." Indeed, the set of models of $\Phi_{>_X}$ is in one-to-one correspondence with the (not-yet-proven-to-be-nonempty) set of total strict orders on X that extend $>_X$, where a model for $\Phi_{>_X}$ is mapped to the order $>'$ defined such that for every distinct $a, b \in X$, we have that $a >' b$ if and only if the formula '${}_a\mathtt{gt}_b$' is TRUE in that model. Thus, Theorem 2.2 is equivalent to $\Phi_{>_X}$ being satisfiable for any given X and strict partial order $>_X$ over the elements of X, and by Compactness it is enough to show that every finite subset of the $\Phi_{>_X}$ is satisfiable.

Thus Theorem 2.2 follows immediately from Theorem 2.1: Every finite subset $\Phi' \subset \Phi_{>_X}$ "mentions" only finitely many elements of X; we denote the set of these elements by $X' \subset X$. It is then immediate that Φ' is satisfiable, as by the finiteness of X' there is *some* strict total order over X' that extends the strict partial order $>_X|_{X'}$ (e.g., Lahiri, 2002); the model corresponding to that order satisfies Φ'.[16] As we have shown that any finite subset of $\Phi_{>_X}$ is satisfiable, we know from Theorem 2.1 that $\Phi_{>_X}$ is satisfiable as well. $\square$

[15] Mas-Colell, Whinston, and Green (1995) label their proof (which uses Zorn's Lemma) as "advanced." Similarly, Ok (2007, p. 17) explains that "[a]lthough it is possible to prove this [fundamental result of order theory] by mathematical induction when X is finite, the proof in the general case is built on a relatively advanced method[. . .]."

[16] Such a model is in fact a model for all formulae in $\Phi_{>_X}$ that "mention" only elements from X', of which Φ' is a subset, and so this is a model for Φ' as well.

Throughout the paper, we apply Theorem 2.1 to show the existence of a solution to a variety of infinite-size or infinite-time economic problems in a way that roughly follows the outline of the proof of Theorem 2.2 just presented: for each problem, we show how to construct a set of well-formed formulae that corresponds to the infinite problem in the sense that any model that satisfies that set "encodes" a solution to the infinite problem, and such that any finite subset corresponds (in the same sense) to a finite variant of the infinite problem. Thus, known results for finite variants of our problems imply models for any finite subset of the formulae, and so, by the Compactness Theorem, we obtain the existence of a solution to the infinite problem. As already hinted in the preceding discussion, one challenge in formulating a set of formulae that corresponds to an infinite problem is to do so in such a way that each of the formulae really is, by itself, finite.

3 Warm Up: Existence of Stable Matching in Infinite Markets

3.1 Setting

We work with the simplest possible matching market setting: a one-to-one "marriage" matching market. Such a market is represented by a quadruplet $(M, W, \mathcal{P}_M, \mathcal{P}_W)$, where M is a (possibly infinite[17]) set of men, W is a (possibly infinite) set of women, and $\mathcal{P}_M$ is a profile of preferences for the men over the women consisting, for each man $m \in M$, of a linearly ordered preference list of women that either is finite, or specifies man m's nth-choice woman for every $n \in \mathbb{N}$. Any woman on m's list is considered preferred by m over being unmatched, while any woman not on m's list is considered unacceptable to m. Similarly, $\mathcal{P}_W$ is a profile of preferences for the women over the men. A (one-to-one, not necessarily perfect) *matching* between M and W is a pairwise-disjoint set of man-woman pairs. A *blocking pair* with respect to a matching μ is a man-woman pair (m, w) such that m prefers w to his partner in μ (or, if he is unmatched in μ, prefers w to being unmatched) and w prefers m to her partner in μ (or, if she is unmatched in μ, prefers m to being unmatched).

A matching μ is called *stable* if (1) under μ, no participant is matched to a partner he or she finds unacceptable (individual rationality), and (2) there are no blocking pairs with respect to μ. Of note, if an agent does not have a best choice from some menu of potential partners, then a stable matching need not exist.

3.2 Existence

As a warm-up, we use our approach to give a simple (re-)proof of a known result on the existence of stable matchings in infinite, one-to-one matching markets.

A classic result of Gale and Shapley (1962) shows that stable matchings exist for any *finite* matching market in the setting just described.

Theorem 3.1 (Gale and Shapley, 1962)**.** *In any finite, one-to-one matching market, a stable matching exists.*

Our Logical Compactness approach gives us a way to lift Theorem 3.1 to infinite markets. (For alternative proofs via a fixed-point argument, or—for a special case—via an infinite variant of Gale and Shapley's algorithm, see Fleiner (2003) and Jagadeesan (2018b), respectively.)

Theorem 3.2. *In any (possibly infinite) matching market, a stable matching exists.*

[17] While (as we describe soon) we must require that each agent finds at most countably many agents acceptable, we make no assumptions on the cardinality of the set of agents.

Proof. We start by defining a set of variables (atomic formulae) $V_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$ and a set of formulae $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$ over those variables, such that the possible models that satisfy $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$ are in one-to-one correspondence with stable matchings in $(M, W, \mathcal{P}_M, \mathcal{P}_W)$. As already mentioned in the Introduction, for this proof we will have for every man $m \in M$ and woman $w \in W$ a variable $\texttt{matched}_{(m,w)}$ that will be TRUE in a model if and only if m and w are matched (in the matching corresponding to the model). So, we will have

$$V_{(M,W,\mathcal{P}_M,\mathcal{P}_W)} \triangleq \{\texttt{matched}_{(m,w)} \mid m \in M \;\;\&\;\; w \in W\}.$$

We now proceed to define the set of formulae $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$.

- For every man m and for every two women $w \neq w'$, we add the following formula to $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$:
$$\texttt{matched}_{(m,w)} \rightarrow \neg\texttt{matched}_{(m,w')}, \tag{5}$$
requiring that m be matched to at most one woman.

- For every woman w and for every two men $m \neq m'$, we add the following formula to $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$:
$$\texttt{matched}_{(m,w)} \rightarrow \neg\texttt{matched}_{(m',w)}, \tag{6}$$
requiring that w be matched to at most one man.

- For every man m and woman w such that either m finds w unacceptable or w finds m unacceptable, we add the following formula to $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$:
$$\neg\texttt{matched}_{(m,w)}, \tag{7}$$
requiring that no one is matched to someone that he or she finds unacceptable.

- For every man m and woman w such that neither finds the other unacceptable, let $w_1, \ldots, w_l$ be all the women that m prefers to w and let $m_1, \ldots, m_k$ be all the men that w prefers to m. (Note that l and k are finite even if the preference lists of w or m are infinite.[18]) We add the following (finite!) formula to $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$:
$$\neg\texttt{matched}_{(m,w)} \rightarrow \big((\texttt{matched}_{(m,w_1)} \vee \cdots \vee \texttt{matched}_{(m,w_l)})\vee \\ \vee (\texttt{matched}_{(m_1,w)} \vee \cdots \vee \texttt{matched}_{(m_k,w)})\big), \tag{8}$$
requiring that (m, w) is not a blocking pair. (Recall that by definition, for this formula to hold either the left-hand side must be FALSE, i.e., m and w must be matched, or the right-hand side must be TRUE, i.e., one of m and w must not prefer the other to her match.)

By construction, and by definition, the models that satisfy all the formulae specified in (5) and (6) are in one-to-one correspondence with matchings between M and W. Furthermore, the models that satisfy $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$ (i.e., all the formulae specified in (5)–(8)) are in one-to-one correspondence with stable matchings between M and W. As noted above, the crux of our argument is that we were able to formalize a set of (*individually finite*) formulae with this property. So, it is

[18] Our assumption of a preference list being of the order type of the natural numbers is precisely what allows us to express stability via individually finite formulae—as required for the Compactness Theorem to be applicable. Fleiner (2003) studies a model with infinite preference lists of more general order types (beyond which stable matchings are known not to exist), under which such a construction would not be possible.

enough to show that $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$ is satisfiable, and by the Compactness Theorem, it is enough to show that every finite subset $\Phi' \subseteq \Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$ is satisfiable.

Let Φ' be a finite subset of $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$. Let M' be the set of all men m such that a variable $\texttt{matched}_{(m,w)}$, for some $w \in W$, appears in one or more formulae in Φ'. Let W' be the set of all women w such that a variable $\texttt{matched}_{(m,w)}$, for some $w \in W$, appears in one or more formulae in Φ'. Since Φ' is finite, only finitely many variables appear in the formulae in Φ', and hence both M' and W' are finite. Let $\mathcal{P}'_{M'}$ be the preferences of M' (induced by $\mathcal{P}_M$), restricted to W', and let $\mathcal{P}'_{W'}$ be the preferences of W' (induced by $\mathcal{P}_W$), restricted to M'. Since $\Phi' \subseteq \Phi_{(M',W',\mathcal{P}'_{M'},\mathcal{P}'_{W'})}$ (which we define analogously for that market), every model that satisfies the latter also satisfies the former. By Theorem 3.1, the latter is satisfiable. Therefore, so is the former, and therefore, by the Compactness Theorem, so is $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$. Therefore, a model that satisfies $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$ exists, and so a stable matching exists for $(M, W, \mathcal{P}_M, \mathcal{P}_W)$. $\square$

While one can prove Theorem 3.2 using other methods (again, see Fleiner, 2003; Jagadeesan, 2018b), as we show in Section 4, Compactness lets us extend structural and strategic results for matching, as well. Additionally, in Section 4 we demonstrate how to utilize the same Compactness argument used in this section, with minimal changes, to prove existence results for variant settings, even those in which the finite case has been analyzed using completely different tools.

3.3 A Cautionary Tale of a Non-Proof

As described in the Introduction, once we were able to translate the stable matching problem into a set of locally finite logical statements (here, statements about variables specifying who matches with whom), Logical Compactness let us take the classical existence result for finite markets (Theorem 3.1) and lift it to infinite markets (Theorem 3.2). One has to be careful, though, beyond making sure that each logical formula is finite. Consider, for example, the following "alternative" to all formulae specified in (8) in the preceding proof:

- For every two men m, m' and for every two women w, w' such that m prefers w over w' and w prefers m over m', add the following formula:

$$\neg\big(\texttt{matched}_{(m,w')} \wedge \texttt{matched}_{(m',w)}\big). \tag{8'}$$

Formula (8') seems to also preclude the possibility of a blocking pair (namely, m and w), and is certainly simpler than Formula (8). Nonetheless, upon closer inspection, we see that there is a trivial model that satisfies Formulae (5)–(7) as well as Formula (8'): the model where $\texttt{matched}_{(m,w)}$ is FALSE for *every* m and w. So, an existence of a model for this "alternative" set of formulae (with Formula (8')) does not guarantee the existence of a stable matching. Indeed, we have crafted the original Formula (8) to also remove the possibility of this model. One may be tempted to fix this for example by introducing a formula that says that each participant must be matched, or alternatively adding formulae analogous to (8') for the case where, say, m is unmatched (rather than matched to w'), however any such solution requires us to be able to express the concept "m is unmatched" in our formulae, which requires an infinite disjunction that cannot be expressed via finite formulae.[19]

[19] Any attempt to circumvent this concern by adding a variable that is TRUE if and only if m is unmatched similarly fails, as forcing a variable to be TRUE if all $\texttt{matched}_{(m,w)}$ are FALSE again requires an infinite disjunction.

3.4 Remarks on Limitations

In the next section we will set out to prove the strategy-proofness of the man-optimal stable matching mechanism. We previously mentioned that our proof of this result will not rely on the Lone Wolf/Rural Hospitals Theorem,[20] while the latter was shown by Jagadeesan (2018b) via a counter-example to not extend to infinite settings. A natural question is therefore where would our approach have failed had we attempted to use it to lift the proof of this latter theorem. The answer is that a Propositional Logic formulation does not allow the free use of quantifiers (e.g., "*there exists* a man m to whom woman w is matched") and similarly does not allow the use of infinitely long expressions such as infinite disjunctions (e.g., "woman w is matched to *one of the following infinitely many* men"). Furthermore, it does not allow statements that condition truths in one model upon truths in other models, such as requiring that a certain statement (e.g., that a certain agent is matched) hold either in all models of our formulae or in none of them. All of these, taken together, preclude capturing the the Lone Wolf/Rural Hospitals Theorem using a Propositional Logic formulation.

As another example, consider the following question:[21] Can our approach be used to show that there exists a mechanism that is strategy-proof, efficient, and individually rational, in a housing market when the sets of objects and agents are infinite, à la Shapley and Scarf (1974)? The answer is that such a construction is at least not straightforward because it is unclear whether and how (Pareto) efficiency can be expressed using (possibly infinitely many) individually finite formulae. Our approach can certainly "lift" to infinite settings the existence of a mechanism that is strategy-proof, individually rational, and rules out any finite trading cycles. However, unlike in the finite setting, in an infinite setting the absence of finite trading cycles does not imply efficiency, as there may exist Pareto improvements that require infinitely many trades. To use our approach, one would need to come up with individually finite formulae (potentially infinitely many of them) that, taken together, rule out inefficiency. Similarly, subsequent to our work, Choi (2020) used our approach to lift group strategy-proofness of various mechanisms into an infinite setting; however, her result only guarantees the lack of incentive for any *finite* coalition to misreport. Extending to infinite coalitions using the same proof technique would once again require a way to reason about deviations by infinite coalitions using (potentially infinitely many) individually *finite* formulae, which seems unlikely to be possible.

4 Infinite Markets: Stable Matchings and Walrasian Equilibria

4.1 Stable Matching: Structure and Incentives

In this section we dive deeper into the two-sided matching setting of Section 3, proving structural and incentive results. The incentive result is novel, and in fact answers a standing open question. Our proofs here are more involved conceptually than those given in the previous sections in several senses: they require working both with the logical model and with the original "semantic" matching model at the same time, restricting to more carefully chosen finite markets and constructing more carefully—and arguably less intuitively—the formulae that define the solution, and also invoking additional results from the literature on finite matching markets. These proofs together also show-

[20] The Lone Wolf Theorem states that in a finite one-to-one matching market, each participant is either matched in all stable matchings, or unmatched in all stable matchings. The Rural Hospitals Theorem is an extension of this theorem for many-to-one matching markets.

[21] We thank an anonymous referee for the 21st ACM Conference on Economics and Computation (EC 2020) for raising this question.

case the potential to use results from the finite setting *beyond the result being lifted* to craft more complex Compactness-based arguments.

4.1.1 Structure

In contrast to our nonconstructive proof of Theorem 3.2, Gale and Shapley's proof of Theorem 3.1 is by way of a constructive argument spelling out an algorithm for finding a stable matching. Tracing the execution of Gale and Shapley's argument gives rise to structural insights about the set of stable matchings, such as the second main result of Gale and Shapley (1962): the existence of a *man-optimal stable matching*, that is, a stable matching that is most preferred (among all stable matchings) by all men simultaneously.

Theorem 4.1 (Gale and Shapley, 1962)**.** *In any finite, one-to-one matching market, there exists a man-optimal stable matching.*

Given the non-algorithmic nature of our proof of Theorem 3.2, it is not *a priori* obvious that the same approach can be used to lift Theorem 4.1 to infinite markets. Indeed, while Theorem 4.1 is also an existence result, it is a far more intricate one, which may be thought of as one that involves an additional level of quantification: the properties of the stable matching whose existence it proves are phrased in terms of all other stable matchings, whose existence we proved in Theorem 3.2. Nonetheless, we can lift Theorem 4.1 to infinite markets by way of Logical Compactness, as well—but the argument is far more intricate than our proof of Theorem 3.2.[22]

Theorem 4.2. *In any (possibly infinite) one-to-one matching market, there exists a man-optimal stable matching.*

As already noted above, the proofs that we present in this section are in some sense conceptually even more involved than those presented so far, and this already manifests in the proof of Theorem 4.2, which is the relatively simpler of the two:

- The argument requires working both with the logical model and with the original "semantic" matching model at the same time in a far more intimate way;
- the argument restricts to more carefully chosen finite markets—and not to "the market of all participants mentioned in any formula in the given finite subset of formulae";
- the formulae constructed may not be the most straightforward way to model the object whose existence we wish to show—they are chosen in order to tweak the problem of satisfiability of a finite subset to allow us to use additional results from the literature on finite matching markets in critical steps of the proof of the satisfiability of such a subset.

Proof of Theorem 4.2. From Theorem 3.2 we know that a stable matching exists. Let $\widehat{M}$ be the subset of men who are matched in at least one stable matching. For each $m \in \widehat{M}$, let w^m be the woman most preferred by m of all women to which he is matched in at least one stable matching.

[22] Again, like with Theorem 3.2, Theorem 4.2 was originally proven by Fleiner (2003). Fleiner's result is in fact more general, and proves the existence of a lattice structure on the set of stable matchings. While it is unclear how to reprove that result using our approach, later in this section we prove a result—namely, the strategy-proofness of the man-optimal stable matching mechanism—that it is unclear how to prove using Fleiner's approach, demonstrating how Fleiner's approach and ours complement each other.

We continue working with the same set of variables $V_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$ as in the proof of Theorem 3.2, however as the set of formulae Φ we take the formulae $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$ from that proof, and for each $m \in \widehat{M}$, also add the following (finite!) formula:

$$\bigvee_{w \succcurlyeq_m w^m} \texttt{matched}_{(m,w)},$$

requiring that m be matched to a woman he prefers at least as much as the woman to whom he is matched in his most-preferred stable matching (of the entire market).[23]

By the definitions of $\widehat{M}$ and w^m, we know that had we added only one such formula (for any choice of one man $m \in \widehat{M}$) to the set of formulae $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$ from the proof of Theorem 3.2, the resulting set of formulae would still by satisfied by some model (corresponding to a stable matching in which m is matched with w^m, which exists by definition of w^m). To prove Theorem 4.2 we would like to show that the set of formulae Φ that includes all of the above-mentioned formulae is satisfied by some model. This will imply the existence of a stable matching in which each man is matched to a woman he prefers at least as much as the woman most preferred by him of all women to which he is matched in any stable matching, hence the existence of a man-optimal stable matching.

We proceed via Logical Compactness. Consider a finite subset of the set of formulae Φ. Since the subset is finite, and each formula is finite, only finitely many men and women are "mentioned" in formulae in the subset. Denote the set of "mentioned" men by M' and the set of "mentioned" women by W'. If the finite subset of formulae does not include any of the new formulae (the ones added in this proof), we know from Theorem 3.1 that this subset is satisfied by some model. Therefore, we assume henceforth that the finite subset of formulae includes at least one of the new formulae.

We know that for every $m \in M' \cap \widehat{M}$ there exists a stable matching of the entire market, μ^m, such that m is matched to w^m. Consider the finite economy consisting of M', W', and the women $\bigcup_{m' \in M' \cap \widehat{M}} \mu^{m'}(M')$. Since this economy is finite, by Theorem 4.1 it has a man-optimal stable matching—we will show that this matching satisfies all of the formulae in the finite subset of formulae.[24]

For any $m \in M' \cap \widehat{M}$, the set of women in this finite economy is a superset of $W' \cup \mu^m(M')$ and the set of men in this finite economy is a subset of $M' \cup \mu^m(W')$. Theorem 2.25 in Roth and Sotomayor (1990) thus assures that the man-optimal stable matching in this finite economy is weakly preferred by all men in M', and in particular weakly preferred by m, to any stable matching in the finite economy consisting of M', W', $\mu^m(M')$ and $\mu^m(W')$. But one of the stable matchings in the latter economy is the restriction of μ^m to this economy (it is stable since any blocking pair would also block the matching μ^m in the full economy), which matches m with w^m. Thus, the man-optimal stable matching of the finite economy consisting of M', W', and $\bigcup_{m' \in M' \cap \widehat{M}} \mu^{m'}(W')$ matches m to a woman he weakly prefers to w^m, and this holds for every man $m \in M' \cap \widehat{M}$. Therefore, this stable matching satisfies all the formulae in the finite subset of formulae.

[23] While we could have instead simply added the formula $\texttt{matched}_{(m,w^m)}$, we add this seemingly more permissive formula. We write "seemingly more permissive" since it in fact would have resulted in the exact same model(s) satisfying the set of formulae Φ! So why do we insist on adding a more elaborate formula if it is in fact not more permissive? We do so because, as we will see soon, this formula may in fact be more permissive *when we consider only a finite subset of* Φ. Indeed, our proof that every such finite subset is satisfiable does not work without using these more permissive formulae. While every subset would be satisfiable either way, using the more permissive formulae causes some finite subsets to be satisfied by more models, and thus an existence of a satisfying model may in fact be easier to prove.

[24] As promised, we have restricted our attention to a finite market far more carefully chosen than simply the market (M', W'). As we will see below, it will indeed be easier for us to reason about stable matchings in this market than in the market (M', W').

By the Compactness Theorem, the entire set of formulae is satisfied by some model, and so a man-optimal stable matching exists in the infinite market. $\square$

4.1.2 Incentives

With Theorem 4.2 in hand, we can define the *man-optimal stable matching mechanism*, which, given any preference profile for the participants in the market, outputs the man-optimal stable matching with respect to those preferences. In finite markets, a classic incentives result of Dubins and Freedman (1981) and Roth (1982) shows that man-optimality leads to *strategy-proofness* (for the men), in the sense that the man-optimal stable matching mechanism makes truthfully reporting preferences a dominant strategy for each man in the market.

Theorem 4.3 (Dubins and Freedman, 1981; Roth, 1982)**.** *In any finite, one-to-one matching market, the mechanism that implements the man-optimal stable matching with respect to reported preferences is strategy-proof for men.*

The challenge in using Logical Compactness to generalize Theorem 4.3 to an infinite setting is threefold. First, as noted already, for finite markets Gale and Shapley (1962) devised an algorithm for finding the man-optimal stable matching—and we can compare the execution of that algorithm under different preference profiles to derive strategy-proofness. In infinite markets, however, while we managed to show the existence of a man-optimal stable matching—and thus that the man-optimal stable matching mechanism is well-defined—we have not given any constructive way to reach it. Second, standard arguments for strategy-proofness rely on versions of the Lone Wolf/Rural Hospitals Theorem (Roth, 1984), which Jagadeesan (2018b) has shown does not hold in our setting (or even in less general infinite-market environments), so a significant innovation on the proof strategy is needed here—and moreover, the mere applicability of Theorem 4.3 to infinite markets is not *a priori* clear. Third, Theorem 4.3 is of a very different flavor than results that have traditionally been lifted to infinite settings using Logical Compactness. Indeed, it is not a result on the existence of a stable matching, and moreover, it does not describe any structural property of a stable matching, but rather an elusive game-theoretic/economic property of a function from preference profiles to stable matchings. Colloquially, if we said that Theorem 4.2 may be thought of as involving an additional level of quantification (there exists a matching such that for every stable matching. . .), then Theorem 4.3 builds on top of it by introducing yet another level of quantification (for every deviation. . .), and furthermore seems quite far from a standard existence result—and what Logical Compactness most naturally helps us prove is existence. Nevertheless, as we show, Logical Compactness, when used in just the right way, lets us prove an existence claim to which the strategy-proofness of the man-optimal stable mechanism can be carefully reduced, and hence to prove that the mechanism itself is strategy-proof.

Theorem 4.4. *In any (possibly infinite) matching market, the mechanism that implements the man-optimal stable matching with respect to reported preferences is strategy-proof for men.*

The proof of Theorem 4.4 has all of the complexities described previously before the proof of Theorem 4.2, along with one more:

- We have to find a way to reduce Theorem 4.4—which is quite far from an existence result—to a result to which we can apply the Compactness Theorem.

As we show, the trick is to focus on the matchings that arise under candidate manipulations of the man-optimal stable matching mechanism.

Proof of Theorem 4.4. We will show that if any man $\tilde{m} \in M$ can manipulate to get matched to any woman $\tilde{w} \in W$, then by reporting truthfully $\tilde{m}$ is matched to a woman he prefers at least as much as $\tilde{w}$. So, let $\mathcal{P}_M = (P_m)_{m \in M}$ and $\mathcal{P}_W$ be preference profiles, let $\tilde{m}$ be a man, and let $P'_{\tilde{m}}$ be alternative preferences for $\tilde{m}$ to report in lieu of $P_{\tilde{m}}$. Let $\tilde{w}$ be the woman matched to $\tilde{m}$ as a result of this manipulation.

We first claim that by reporting to the mechanism the preference list $\{\tilde{w}\}$ consisting solely of $\tilde{w}$ rather than the preference list $P'_{\tilde{m}}$ (when all other agents' reported preferences are still fixed at $\mathcal{P}_{M \setminus \{\tilde{m}\}}$ and $\mathcal{P}_W$), man $\tilde{m}$ will still be matched to $\tilde{w}$. To see this, recall that $\tilde{m}$ is matched with $\tilde{w}$ under the man-optimal stable matching with respect to the manipulated report $P'_{\tilde{m}}$ (and the profile of all others' preferences). But by erasing all other potential spouses (other than $\tilde{w}$) from $\tilde{m}$'s preference report, stability is not compromised, as there are only fewer potential blocking pairs. Hence, $\tilde{w}$ is matched to $\tilde{m}$ in some stable matching with respect to the reported preference list for $\tilde{m}$ that consists solely of $\tilde{w}$, and since she is the only woman on $\tilde{m}$'s preference report, she must be matched to him under the man-optimal stable matching with respect to this list. Let μ be the man-optimal stable matching in the market $(M, W, (\{\tilde{w}\}, \mathcal{P}_{M \setminus \{\tilde{m}\}}), \mathcal{P}_W)$.

Recall that we aim to prove that $\tilde{m}$ will also be matched with $\tilde{w}$ or a (truly) better-preferred woman by reporting his true preferences. For this, we use Logical Compactness.[25]

We work with the variables $V := V_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$ as defined in the proof of Theorem 3.2, and as the set of formulae Φ we take the formulae $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$ from that proof and add to them the following (finite!) formula:

$$\bigvee_{w \succeq_{P_{\tilde{m}}} \tilde{w}} \texttt{matched}_{(\tilde{m},w)},$$

requiring that $\tilde{m}$ be matched to a women he truly prefers at least as much as $\tilde{w}$.

Given a finite subset of these formulae, only a finite set of men and women are mentioned. We call these sets M' and W', respectively, and we henceforth consider the finite market consisting of the men $M' \cup \{\tilde{m}\} \cup \mu(W')$ and the women $W' \cup \{\tilde{w}\} \cup \mu(M')$. We will claim that the man-optimal stable matching in this finite market with respect to the induced profile of preferences when $\tilde{m}$'s preferences are his true preferences $P_{\tilde{m}}$ satisfies the finite subset of formulae.

We know that the same finite market has a stable matching, with respect to the induced profile of preferences when $\tilde{m}$'s list consists solely of $\tilde{w}$, that matches $\tilde{m}$ to $\tilde{w}$ (the restriction of μ to this market). This means that the man-optimal stable matching in this finite market when $\tilde{m}$'s list consists solely of $\tilde{w}$ matches $\tilde{m}$ to $\tilde{w}$. So, by Theorem 4.3, it must be that in the same finite market, the man-optimal stable matching with respect to $P_{\tilde{m}}$ has $\tilde{m}$ matched to a woman he $P_{\tilde{m}}$-ranks at least as high as $\tilde{w}$ (otherwise, had $\tilde{m}$'s true preferences been $P_{\tilde{m}}$, he would have had a profitable manipulation: declaring his list to consist solely of $\tilde{w}$). Since the finite set of formulae can, at most, require stability with respect to mentioned individuals and that $\tilde{m}$ be matched to a woman he weakly prefers to $\tilde{w}$, the above argument establishes that the finite set of formulae is satisfied by some model (corresponding to the man-optimal stable matching in that finite market with respect to $P_{\tilde{m}}$). Hence, by the Compactness Theorem the entire collection of formulae is satisfied by some model, and hence this model corresponds to a stable matching with respect to the true profile of preferences $(\mathcal{P}_M, \mathcal{P}_W)$, in which $\tilde{m}$ is matched to $\tilde{w}$ or a woman he prefers more, and hence in the man-optimal stable matching with respect to the true preferences $\tilde{m}$ is matched to a woman he prefers at least as much as $\tilde{w}$, which is what we set out to prove. $\square$

We emphasize that unlike Theorems 3.2 and 4.2, Theorem 4.4 is novel to the present work—

[25] As promised, we will recast this statement as an existence problem—and it is to this existence problem to which we reduce strategy-proofness, and to which Logical Compactness can be applied.

although a special case in which all preference profiles are finite was proved by Jagadeesan (2018b). Moreover, as already mentioned, unlike Theorems 3.2 and 4.2, Theorem 4.4 is not an existence result about matchings with certain properties—it deals with the incentive properties of a specific matching mechanism. Thus, Theorem 4.4 illustrates that Logical Compactness has applications beyond results we might naturally expect to be able to prove with limiting or continuity arguments.

4.2 Variant Settings: Stable Matching with Couples as a Case Study

Having demonstrated the robustness of the lifting argument that we used for one-to-one matching, it may not seem surprising that it also extends to many-to-one and even to many-to-many matchings. The strength of our Compactness approach, though, is that it extends even to similar models for which the finite existence results hinges on very different tools. For example, while Theorem 3.1 and other classical existence results for finite matching markets are all proven using variants of the deferred acceptance algorithm (Gale and Shapley, 1962), a recent existence result for stable matchings with couples (Nguyen and Vohra, 2018) is proven using a completely different argument that builds upon Scarf's Lemma. While this difference in techniques makes it quite unlikely that the techniques that Fleiner (2003) and Jagadeesan (2018b) applied to prove Theorem 3.2 would be applicable to proving an infinite-market analogue of the main result of Nguyen and Vohra (2018), we will show as an example in this section that our approach generalizes readily to lift this result as well.

Like Nguyen and Vohra (2018), we study the standard matching with couples model (e.g., Roth, 1984; Klaus and Klijn, 2005; Kojima et al., 2013; Ashlagi et al., 2014). In this model, a market is a tuple $(D, H, \mathcal{P}_D, \mathcal{P}_H, k)$ where D is the union of the set D^1 and the set D^2 of couples, H is a set of hospitals, $k \equiv \{k_h\}_{h \in H}$ is a vector of hospital capacities, $\mathcal{P}_H$ is a profile of rankings for the hospitals over the doctors consisting, for each hospital $h \in H$, of a linearly ordered ranking over doctors that either is finite, or specifies hospital h's nth-choice doctor for every $n \in \mathbb{N}$. Hospitals preferences are *responsive*—from any set of available doctors they choose the highest-ranked ones up to the hospital's capacity (always rejecting unranked doctors). $\mathcal{P}_D = (\mathcal{P}_{D^1}, \mathcal{P}_{D^2})$, is a profile of doctor preferences. Single doctors' preferences, $\mathcal{P}_{D^1}$, are similarly defined. For couples, the definition is also similar with the exception that $\mathcal{P}_{D^2}$ is a linearly ordered ranking over ordered pairs in $(H \cup \{\emptyset\}) \times (H \cup \{\emptyset\}) \setminus \{(\emptyset, \emptyset)\}$, representing the assignment of the first, and second, member of the couple, that the couple prefer to both being unmatched.

Given a matching-with-couples market and a vector of capacities, k^*, a matching is *individually rational with respect to the capacities* k^* if no single doctor is assigned to an unacceptable hospital, couples are assigned to (h, h') which they weakly prefer to $(h, \emptyset)$, $(\emptyset, h')$, and $(\emptyset, \emptyset)$, and each hospital h is assigned no more than k^*_h doctors, all of whom are ranked by $\mathcal{P}_h$. A matching μ is *blocked with respect to the capacities* k^* if one of the following holds: 1) there exists a single doctor, $d \in D^1$, and a hospital h, such that d prefers h to $\mu(d)$ and h's most preferred subset of $\mu(h) \cup d$, subject to k^*_h, includes d. 2) there exists a couple $c \in D^2$ and a hospital h, such that the couple prefers (h, h) to $\mu(c)$ and h would select both members of the couple as above. Or, 3) there is a triple $(c, h, h') \in D^2 \times (H \cup \{\emptyset\}) \times (H \cup \{\emptyset\})$ with $h \neq h'$ such that the couple prefers (h, h') to $\mu(c)$ and each of the hospitals chooses the respective member of the couple from the set as above. Finally, a matching is *stable with respect to the capacities* k^* if it is individually rational and not blocked.

Theorem 4.5 (Nguyen and Vohra, 2018)**.** *In any finite, many-to-one matching market with couples with capacity vector k, there exists a capacity vector k^* with $|k_h - k^*_h| \leq 2$ for every $h \in H$ and $\sum_{h \in H} k_h \leq \sum_{h \in H} k^*_h \leq \sum_{h \in H} k_h + 4$ such that a stable matching w.r.t. k^* exists.*

Despite the very different proofs of Theorems 3.1 and 4.5 (the former uses deferred acceptance;

the latter uses Scarf's Lemma), our proofs of their infinite extensions are remarkably similar. As in the one-to-one stable case, we will now want to capture the solution concept of stability with couples and approximately same capacities via a set of locally finite logical formulae and then use Logical Compactness to take the existence result for finite markets (Theorem 4.5) and lift it to infinite markets. To make sure that each logical formula that we come up with is indeed finite, we will need a technical assumption on the preferences of couples in the market.

Definition 4.6 (Downward Closed)**.** We say that a preference order for a couple over $(H \cup \{\emptyset\})^2 \setminus \{(\emptyset, \emptyset)\}$ is *downward closed* if for every pair of (actual, i.e., not $\emptyset$) hospitals $(h, h') \in H^2$ ranked by this order, both $(h, \emptyset)$ and $(\emptyset, h')$ are also ranked by this order.

Theorem 4.7. *In any (possibly infinite) matching market with couples with capacity vector k, if the preference of each couple is downward closed,*[26] *then there exists a capacity vector k^* with $|k_h - k^*_h| \leq 2$ for every $h \in H$ such that a stable matching w.r.t. k^* exists.*

The formulae that we build to prove Theorem 4.7 are conceptually similar to those used in our proof of Theorem 3.2, yet somewhat more intricate, due to the many-to-one nature of the market and the presence of couples, as well as to the variability in capacity. The idea is to have, as before, for every doctor $d \in D$ and hospital $h \in H$ a variable $\texttt{matched}_{(d,h)}$ that will be TRUE in a model if and only if d and h are matched in the matching corresponding to the model. Furthermore, for every hospital $h \in H$ with capacity k_h, we have five variables $\texttt{capacity}_{(h,k_h-2)}, \texttt{capacity}_{(h,k_h-1)}, \ldots, \texttt{capacity}_{(h,k_h+2)}$ such that $\texttt{capacity}_{(h,q)}$ will be TRUE in a model if and only if $k^* = q$, and upon whose value each formula will be conditioned—so for instance, for each $q \in \{k_h{-}2, \ldots, k_h{+}2\}$ and for every $q{+}1$ doctors we will have a formula that says "if the capacity of h is q, then h is not matched to these $q{+}1$ doctors." Except for the potential need to perturb the capacities, and the need to express couples' preferences and the absence of blocks involving couples, the main ideas from the proof of Theorem 3.2 carry over; we relegate the details to Appendix B.

4.3 Walrasian Equilibria in Infinite Trading Networks

We next turn to an infinite variant of the trading-network framework of Hatfield et al. (2013) and show the existence of Walrasian equilibria. The application of Compactness here is of a slightly different flavor: we first use the Compactness Theorem to show the existence of arbitrarily close approximations of Nash equilibria in the infinite trading network, and then show that the existence of approximate Nash equilibria implies the existence of exact Nash equilibria.[27]

There is a (potentially infinite) set I of agents. A *trade* ω transfers an underlying object, $\mathsf{o}(\omega)$, from a seller $\mathsf{s}(\omega)$ to a buyer $\mathsf{b}(\omega)$. We denote the set of potential trades by Ω. For $i \in I$ we denote by Ω_i the set of trades in which i participates, namely, $\Omega_i := \{\omega \in \Omega \mid i \in \{\mathsf{s}(\omega), \mathsf{b}(\omega)\}\}$. We assume that Ω_i is finite for every i—that is, each agent is a party to finitely many (potential) trades (note that this implies that each agent is endowed with at most finitely many objects to trade, and has at most finitely many trading partners).

Each agent's utility depends only on the trades that she executes, and the prices at which these trades are executed. Specifically, each agent i is associated with a utility function that is quasilinear in prices and otherwise depends only on the set of trades $\Omega'_i \subseteq \Omega_i$ that are executed.

[26] Alternatively, we can replace the requirement that preferences of couples are downward closed with a requirement that preference lists of couples are finite, and essentially the same proof would go through.

[27] We demonstrate the generality of this use of Logical Compactness by using a similar argument in Appendix C to reprove the existence of a Nash equilibrium in graphical games—a setting for which the finite proof is decidedly different than the proof of the finite proof in the trading-network setting studied in this section.

The "trades" terminology that we use highlights that in our model, objects are linked to specific trading partners—so "car sold to Alice" is a different object than "car sold to Bob," even though the physical good that is traded in reality might be the same car. (To rule out the possibility that the same car is traded to multiple people, the agent's utility function can assign value $-\infty$ to executing "car sold to Alice" and "car sold to Bob" simultaneously.) Having clarified this issue, it will be easier to think about our model in terms of objects from this point on.

We denote by O the set of all objects. Note that objects and trades are in one-to-one correspondence. For an object o, we let $\mathsf{t}(o)$ denote the trade associated with that object, so for each object o we have $\mathsf{o}(\mathsf{t}(o)) = o$, and for each trade ω we have $\mathsf{t}(\mathsf{o}(\omega)) = \omega$. For each agent i we denote by $O_i := \{o \in O \mid i \in \{\mathsf{s}(\mathsf{t}(o)), \mathsf{b}(\mathsf{t}(o))\}\}$ the set of all objects that can be held by i (and recall that $|O_i| = |\Omega_i|$ is finite by assumption). Each $o \in O$ belongs to exactly two sets in $\{O_i\}_{i \in I}$. We may think of i's utility function as expressing the value of the objects that i "holds" after the execution of trades (that is, the set of objects $\{\mathsf{o}(\omega) \mid \omega \in \Omega'_i \text{ and } \mathsf{b}(\omega) = i\} \cup \{\mathsf{o}(\omega) \mid \omega \notin \Omega'_i \text{ and } \mathsf{s}(\omega) = i\}$; see Hatfield et al., 2019).

We assume that for each i, the utility function $u_i(\cdot) : 2^{O_i} \to \mathbb{R} \cup \{-\infty\}$ takes values in $\mathbb{R} \cup \{-\infty\}$ (again, where we use $-\infty$ to model technological impossibilities such as selling the same car to multiple buyers). We further assume that in the absence of trade i's utility is equal to 0 (formally, $u_i(\{\mathsf{o}(\omega) \mid \mathsf{s}(\omega) = i\}) = 0$; this is a normalization, except in that it rules out some agents "having to" execute certain trades.

For each agent, i, let the *demand correspondence* $D_i : p \in \mathbb{R}^O \rightrightarrows 2^{O_i}$ be the correspondence that is defined by the arg max of agent i's utility under the prices p. The preferences of agent i are *(gross) substitutable* if for all price vectors $p, p' \in \mathbb{R}^O$ such that $|D_i(p)| = |D_i(p')| = 1$ and $p \le p'$, if $o \in D_i(p)$ then $o \in D_i(p')$ for each $o \in O_i$ such that $p_o = p'_o$.

A *Walrasian equilibrium* consists of a vector of prices $p \in \mathbb{R}^O$, and a partition $\{O'_i\}_{i \in I}$ of O, such that $O'_i \in D_i(p)$ for each $i \in I$. Hatfield et al. (2013) have shown that substitutable preferences suffice to guarantee the existence of Walrasian equilibria in finite trading networks.

Theorem 4.8 (Hatfield et al., 2013)**.** *If all agents have substitutable preferences, then every finite trading network has a Walrasian equilibrium.*

The main result of this section is a generalization of Theorem 4.8 to infinite trading networks.

Theorem 4.9. *If all agents have substitutable preferences, then every (possibly infinite) trading network has a Walrasian equilibrium.*

As already noted, we prove Theorem 4.9 by first using Logical Compactness to prove the existence of arbitrarily good approximate Walrasian equilibria, and then show that the existence of such approximate Walrasian equilibria implies Theorem 4.9. We start by defining precisely what we mean by approximate Walrasian equilibria.

For every $\varepsilon > 0$, the *approximate demand correspondence* $D_i^\varepsilon : p \in \mathbb{R}^O \rightrightarrows 2^{O_i}$ is defined similarly to the (exact) demand correspondence, except that its range includes bundles from which agent i's utility is at least that of the utility-maximizing bundle minus ε. (Thus, for all i and p, $D_i(p) \subseteq D_i^\varepsilon(p)$.) For a given $\varepsilon > 0$, an *ε-Walrasian equilibrium* is a vector of prices $p \in \mathbb{R}^O$, and a partition $\{O'_i\}_{i \in I}$ of O, such that $O'_i \in D_i^\varepsilon(p)$ for every $i \in I$. For a given vector $(\varepsilon_i)_{i \in I}$, an *$(\varepsilon_i)_{i \in I}$-Walrasian equilibrium* is a vector of prices, $p \in \mathbb{R}^O$, and a partition $\{O'_i\}_{i \in I}$ of O, such that $O'_i \in D_i^{\varepsilon_i}(p)$ for every $i \in I$.

Lemma 4.10. *If all agents have substitutable preferences, then for every $\varepsilon > 0$, every (possibly infinite) trading network has an $(|O_i| \cdot \varepsilon)_{i \in I}$-Walrasian equilibrium.*

Proof. We first note that for every object $o \in O$, there is a positive integer H_o such that i) if the price of o is H_o (and there is no technological impossibility), then $\mathsf{s}(o)$ always wishes to sell o and $\mathsf{b}(o)$ always wishes to not buy o, and ii) if the price of o is $-H_o$ (and there is no technological impossibility), then $\mathsf{s}(o)$ always wishes to hold o and $\mathsf{b}(o)$ always wishes to buy o. Formally, there exists H_o such that for every $i \in \{\mathsf{s}(o), \mathsf{b}(o)\}$ and $O'_i \subseteq O_i$, if $\left|u_i(O'_i \cup \{o\})\right| + \left|u_i(O'_i)\right| < \infty$ then $\left|u_i(O'_i \cup \{o\}) - u_i(O'_i)\right| < H_o$. (To show existence, just consider any upper bound on these expressions and take a greater integer.)

Let $\varepsilon > 0$ and let n be an integer greater than $1/\varepsilon$. For every object $o \in O$ we denote the set of *possible prices* for o by $P_o := \{-H_o, \ldots, -\frac{1}{n}, 0, \frac{1}{n}, \frac{2}{n}, \ldots, H_o\}$. We once again define a set of variables V and a set of formulae Φ over these variables, such that the models that satisfy Φ are in one-to-one correspondence with $(|O_i| \cdot \varepsilon)_{i \in I}$-Walrasian equilibria in the given trading network in which the price of each $o \in O$ is in P_o. We will have the following variables (atomic formulae) in V:

- $\texttt{price}_{(o,p)}$ for every object $o \in O$ and possible price $p \in P_o$.
- $\texttt{consumes}_{(i,o)}$ for every $i \in I$ and $o \in O_i$.

Formulae of the first type will represent the price of each object, and formulae of the second type will represent which agents hold which objects after the execution of trades. Next, we define the set of formulae Φ as follows:

1. For every object $o \in O$, we add the following (finite!) formula:

$$\texttt{price}_{(o,-H_o)} \vee \texttt{price}_{(o,-H_o+\frac{1}{n})} \vee \cdots \vee \texttt{price}_{(o,H_o)},$$

   requiring that o have a price in P_o.

2. For every object $o \in O$ and for every pair of distinct possible prices, $p, p' \in P_o$, we add the following formula:

$$\texttt{price}_{(o,p)} \to \neg\texttt{price}_{(o,p')},$$

   requiring that the price of o be unique.

3. For every object $o \in O$ we add the following formula:

$$\texttt{consumes}_{(\mathsf{b}(\mathsf{t}(o)),o)} \leftrightarrow \neg\texttt{consumes}_{(\mathsf{s}(\mathsf{t}(o)),o)},$$

   requiring that o either be sold or not. Equivalently, it requires that the associated trade $t(o)$ either be executed or not.

4. For each $i \in I$ and for each vector of possible prices $p = (p_o)_{o \in O_i} \in \bigtimes_{o \in O_i} P_o$, we write the (finite) formula:

$$\Big(\bigwedge_{o \in O_i} \texttt{price}_{(o,p_o)}\Big) \to \Big(\bigvee_{X \in D_i^{|O_i|\varepsilon}(p)} \Big(\bigwedge_{x \in X} \texttt{consumes}_{(i,x)} \wedge \bigwedge_{x \in O_i \setminus X} \neg\texttt{consumes}_{(i,x)}\Big)\Big),$$

   requiring that i consumes one (and only one) of her $(|O_i| \cdot \varepsilon)$-utility-maximizing bundles (and not any object outside this bundle).

Each finite subset of Φ mentions only a finite set of objects, corresponding to a finite set of agents. To satisfy a finite collection of formulae, we set the prices of all objects not mentioned to 0; holding the prices of these objects at 0, by Theorem 4.8 there exists a vector of prices for the mentioned objects and a partition of the objects among the agents that constitute a Walrasian equilibrium. For each mentioned object o, by definition of H_0, if the price of o in this equilibrium price vector is not in $[-H_0, H_0]$, then it can be replaced by H_0 (if it is positive) or by $-H_0$ (if it is negative) without changing the allocation, and we would still have a Walrasian equilibrium. Now, rounding the price of each such o to the nearest $p \in P_o$ causes each corresponding agent's consumption bundle to turn from optimal to (in the worst case) $(|O_i| \cdot \varepsilon)$-optimal. Thus, the rounded prices, together with the equilibrium consumption bundles (which belong to the $(|O_i| \cdot \varepsilon)$-demand with respect to the rounded price vector) satisfy the finite collection of formulae. By the Compactness Theorem, the entire set Φ is therefore also satisfiable. Hence, there exists an a vector of prices in $\bigtimes_{o \in O} P_o$ and a partition of O that constitute an $(|O_i| \cdot \varepsilon)_{i \in I}$-Walrasian equilibrium. $\square$

While Lemma 4.10 by itself does not directly guarantee even the existence of an ε-Walrasian equilibrium (indeed, we merely assumed that the number $|O_i|$ of (potential) trades to which each given agent is a side is finite, and did not assume any uniform bound on $|O_i|$ across agents), using a limit argument over the result of Lemma 4.10 nevertheless yields the existence of an exact Walrasian equilibrium.

Proof of Theorem 4.9. Since each agent in the trading network has finitely many neighbors (agents she can trade with), every connected component of the network consists of at most countably many agents. Since it is enough to show the existence of a Walrasian equilibrium in each connected component separately (we use the Axiom of Choice here), let us focus on one connected component. For a diminishing sequence $\varepsilon_n \to 0^+$, by Lemma 4.10 there exists a sequence of $(|O_i| \cdot \varepsilon_n)_{i \in I}$-Walrasian equilibria in the connected component. As the number of objects in the connected component is countable, we can choose a subsequence (a "diagonal subsequence") such that the price of each object converges—and so does each agent's consumption bundle. Since ε-demand correspondences are upper hemicontinuous (i.e., weak inequalities are preserved in the limit) and since each O_i is finite, for each agent the limit of the subsequence of approximately optimal consumption bundles is an (exact) optimal consumption bundle for the limit prices. Furthermore, markets must clear, as for each object, there exists some large enough index after which the object is always traded or never traded on the subsequence of the $(|O_i| \cdot \varepsilon_n)_{i \in I}$-Walrasian equilibria. Hence, the limit of the subsequence is an exact Walrasian equilibrium. In particular, an exact Walrasian equilibrium exists in the infinite trading network. $\square$

5 Infinite Time: Stable Matching with a Doubly-Infinite Horizon

In this section we use Logical Compactness to prove the existence of stable matchings in dynamic, infinite-size, and infinite-horizon markets, generalizing a dynamic stable matching model of Pereyra (2013) in which time is bounded from below. For simplicity, we formulate the dynamic setting with one-to-one matching (in each period). Like all of the other models in this paper, the dynamic model we consider is motivated by an established framework—in this case, the model of teachers-to-schools assignment with tenure constraints introduced by Pereyra (2013). For consistency with our other matching sections, here we speak of the agents as "men" and "women," even though they are really stand-ins for "teachers" and "schools."

5.1 Setting

A dynamic matching market is a tuple $(M, W, \mathcal{P}_M, \mathcal{P}_W, (a_m)_{m\in M}, (d_m)_{m\in M})$, where $(M, W, \mathcal{P}_M, \mathcal{P}_W)$ is a (possibly infinite) matching market as in Section 3, and where for each $m \in M$, we have that $a_m < d_m$ are integer numbers, respectively called the *arrival time* and *departure time* of m. For each $m \in M$, we say that m is *on the market* at all (integral) times $t \in [a_m, d_m)$. (All $w \in W$ are considered to always be on the market.) A *matching chronology* in a dynamic matching market is a mapping from woman-time pairs to men who are on the market at the relevant time, such that each man is matched to at most one woman at any given time. We say that a matching chronology is *stable subject to tenure* if:

- *Men have tenure*: for every time t, every man who is on the market both at time t and at time $t+1$, weakly prefers his match at time $t+1$ to his match at time t.
- The matching is *otherwise stable*: at any time t, there exists no pair of man m and woman w such that man m strictly prefers w to his match at t, and w strictly prefers m to her match at time t who is furthermore not her match at time $t-1$.[28]

Our dynamic setting builds on the model of Pereyra (2013), in which arrival times are required to be nonnegative. If we were to restrict ourselves to nonnegative arrival times (or more generally, to arrival times that have a finite lower bound), then, following Pereyra (2013), a simple iterative application of the man-optimal stable matching mechanism would find a stable-subject-to-tenure matching chronology:

1. As the matching at time 0, use the man-optimal stable matching for all women and all men with arrival time 0.
2. As the matching at time 1, use the man-optimal stable matching for all women and all men who are on the market at time 1, with respect to slightly modified preferences: any man who is matched at time 0 and still on the market at time 1 is promoted (for the purposes of finding the man-optimal stable matching at time 1) to be top-ranked on the preferences of his match at time 0.
3. As the matching at time 2, use the man-optimal stable matching for all women and all men who are on the market at time 2, with respect to slightly modified preferences: any man who is matched at time 1 and still on the market at time 2, is promoted (for the purposes of finding the man-optimal stable matching at time 2) to be first on the preferences of his match at time 1.
4. . . . and so on.

The preceding argument in fact constitutes a full proof of the following.

Theorem 5.1 (Pereyra, 2013)**.** *In any dynamic matching market where all arrival times are nonnegative, a stable-subject-to-tenure matching chronology exists.*

As a referee noted, our analysis from Section 3 suffices to immediately extend Theorem 5.1 to infinite markets (via the same iterative process of finding successive man-optimal stable matchings). In this section, however, we show a different generalization: that of making time, rather than market size, infinite—or more precisely, doubly infinite.

[28]Note that, following Pereyra (2013), we implicitly assume that agents' preferences over partners are consistent over time (unlike in, e.g., the framework of Kadam and Kotowski, 2018). Additionally, again following Pereyra (2013), we enforce stability myopically (unlike in the frameworks of Doval, 2017; Liu, 2018; Ali and Liu, 2019).

5.2 Challenge

As we have just seen, the proof of Theorem 5.1 depends heavily on our ability to identify a "starting" matching that we can adjust/build off of in subsequent time periods. However, the need to assume a fixed start time makes the model less representative of a steady-state.

What if arrival times have no finite lower bound? Due to symmetry considerations, there can be no "reasonable" deterministic variant of the man-optimal stable matching mechanism that reaches a stable-subject-to-tenure marriage chronology:

Example 5.2. Consider a case of one woman w and an infinite set of men m_t such that for each $t \in \mathbb{Z}$, a man m_t has arrival time t and departure time $t+2$. For any profile of preference lists in which no agent finds any other agent unacceptable, there are precisely two stable-subject-to-tenure matchings chronologies:

- All men with even arrival times are matched to w throughout their time on the market; all men with odd arrival times are never matched.
- All men with odd arrival times are matched to w throughout their time on the market; all men with even arrival times are never matched.

Which of the two preceding stable-subject-to-tenure matching chronologies (none more or less "man-optimal" than the other) would a deterministic variant of the man-optimal stable matching mechanism, if one existed, choose in the setting of Example 5.2? The only way to break the symmetry and choose between these two is to give special treatment to some specific time period, such as 0—but it is easy to see that just picking some finite time t, matching m_t with w, and solving forward and (somehow) backward would not work since the removal of even a single man, say with very small (negative) arrival time t', from the market would collapse the two stable matching chronologies starting at $t'+1$ (with $m_{t'+1}$ being matched in any stable matching).

In the absence of a reasonable variant of the man-optimal stable matching mechanism, we need a new way to prove existence when arrival times are unbounded. To resolve this problem, we turn again to Logical Compactness.

5.3 Existence

To prove the existence of a stable-subject-to-tenure matching chronology in the infinite-history model, we require some mild "local finiteness" conditions (but not local boundedness), which we will show if fact cannot be dropped.

Definition 5.3 (Finite Presence)**.** Let $(M, W, \mathcal{P}_M, \mathcal{P}_W, (a_m)_{m\in M}, (d_m)_{m\in M})$ be a dynamic matching market. If for every time t, only finitely many men are on the market both at t and at $t+1$, then we say that the dynamic matching market has *finite presence.*

Theorem 5.4. *In any dynamic matching market (with arbitrary arrival and departure times) that has finite presence, a stable-subject-to-tenure matching chronology exists.*

Before proving Theorem 5.4, we note that the finite presence condition in that theorem cannot be dropped.

Example 5.5. It is straightforward to verify that in a dynamic market with one woman w and countably many men $\{m_t\}_{t\in\mathbb{N}}$, with $a_{m_t} = -t$ and $d_{m_t} = 0$ for every t, no stable-subject-to-tenure matching chronology exists if all participants find all possible partners acceptable.

Proof of Theorem 5.4. Like in our proof of Theorem 3.2, we define a set of variables

$$V_{(M,W,\mathcal{P}_M,\mathcal{P}_W,(a_m)_{m\in M},(d_m)_{m\in M})}$$

and a set of formulae

$$\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W,(a_m)_{m\in M},(d_m)_{m\in M})}$$

over those variables, such that the models that satisfy $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W,(a_m)_{m\in M},(d_m)_{m\in M})}$ are in one-to-one correspondence with the stable-subject-to-tenure matching chronologies in $(M,W,\mathcal{P}_M,\mathcal{P}_W,(a_m)_{m\in M},(d_m)_{m\in M})$. This time, for every man $m\in M$, woman $w\in W$, and time $t\in\mathbb{Z}$ we will have a variable $\texttt{matched}_{(m,w,t)}$ that will be TRUE in a model if and only if m and w are matched at time t (in the matching chronology that corresponds to the model). So, we will have $V_{(M,W,\mathcal{P}_M,\mathcal{P}_W,(a_m)_{m\in M},(d_m)_{m\in M})} \triangleq \{\texttt{matched}_{(m,w,t)} \mid m\in M \;\;\&\;\; w\in W \;\;\&\;\; t\in\mathbb{Z}\}$. We define the set of formulae $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W,(a_m)_{m\in M},(d_m)_{m\in M})}$ as follows.

1. For every man m, time t, and two women $w\neq w'$, we add the following formula to $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W,(a_m)_{m\in M},(d_m)_{m\in M})}$:

   $$\texttt{matched}_{(m,w,t)} \rightarrow \neg\texttt{matched}_{(m,w',t)},$$

   requiring that man m be matched to at most one woman at time t.

2. For every woman w, time t, and two men $m\neq m'$, we add the following formula to $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W,(a_m)_{m\in M},(d_m)_{m\in M})}$:

   $$\texttt{matched}_{(m,w,t)} \rightarrow \neg\texttt{matched}_{(m',w,t)},$$

   requiring that woman w be matched to at most one man at time t.

3. For every woman w, time t, and man m such that either m is not on the market at t or w finds m unacceptable or m finds w unacceptable, we add the following formula to $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W,(a_m)_{m\in M},(d_m)_{m\in M})}$:

   $$\neg\texttt{matched}_{(m,w,t)},$$

   requiring that no one is matched to someone that they find unacceptable, and that men not be matched when they are not on the market.

4. For every time t, every woman w, and every man m who is on the market at time t such that neither finds the other unacceptable, let $w_1,\ldots,w_l$ be all the women that m prefers to w, let $m_1,\ldots,m_k$ be all the men that w prefers to m, and let $m'_1,\ldots,m'_n$ be all the men that are on the market at both t and $t-1$. (Note that l and k are finite even if the preference list of w or m is infinite, and that n is finite by finite presence.) We add the following (finite!) formula to $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W,(a_m)_{m\in M},(d_m)_{m\in M})}$:

   $$\begin{aligned}\neg\texttt{matched}_{(m,w,t)} \rightarrow \Big(& (\texttt{matched}_{(m,w_1,t)} \vee\cdots\vee \texttt{matched}_{(m,w_l,t)})\vee \\ & \vee \big((\texttt{matched}_{(m_1,w,t)} \vee\cdots\vee \texttt{matched}_{(m_k,w,t)}) \wedge \neg\texttt{matched}_{(m,w,t-1)}\big)\vee \\ & \vee \big((\texttt{matched}_{(m'_1,w,t)} \wedge \texttt{matched}_{(m'_1,w,t-1)}) \vee\cdots\vee (\texttt{matched}_{(m'_n,w,t)} \wedge \texttt{matched}_{(m'_n,w,t-1)})\big)\Big),\end{aligned}$$

   requiring that (m,w) not be a blocking pair at time t.

The proof concludes via the Compactness Theorem just as in the proof of Theorem 3.2. In particular, note that a finite subset of $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W,(a_m)_{m\in M},(d_m)_{m\in M})}$ only involves finitely many times, and so to show that this finite subset is satisfiable we can apply Theorem 5.1 using the minimum involved time as a "period 0." $\square$

6 Infinite Data: Revealed Preferences

6.1 Rationalizing Consumer Demand

We now move to a very different context—that of rationalizing the consumption behavior of a single consumer. For the most part, in this section we follow the notation of Reny (2015). Fix a number of goods $m \in \mathbb{N}$. A dataset $S \subset \left(\mathbb{R}^m_+ \setminus \{\bar{0}\}\right) \times \mathbb{R}^m_+$ with generic element $(\bar{p}, \bar{x}) \in S$ represents a set of observations, wherein each, a consumer with a budget faces a price vector $\bar{p} \neq \bar{0}$ and chooses to consume the bundle $\bar{x}$. A utility function $u : \mathbb{R}^m_+ \to \mathbb{R}$ *rationalizes* the dataset S if for every $(\bar{p}, \bar{x}) \in S$ and every $\bar{y} \in \mathbb{R}^m_+$, it holds that if $\bar{p} \cdot \bar{y} \leq \bar{p} \cdot \bar{x}$ (i.e., $\bar{y}$ can also be bought with the budget) then $u(\bar{y}) \leq u(\bar{x})$, and if $\bar{p} \cdot \bar{y} < \bar{p} \cdot \bar{x}$ (i.e., $\bar{y}$ can be bought without spending the entire budget) then $u(\bar{y}) < u(\bar{x})$.[29] If only the former implication holds for every such $(\bar{p}, \bar{x})$ and $\bar{y}$, then we will say that u *weakly rationalizes* S.

A dataset S satisfies the *Generalized Axiom of Revealed Preference (GARP)* if for every (finite) sequence $(\bar{p}_1, \bar{x}_1), \ldots, (\bar{p}_k, \bar{x}_k) \in S$, if for every $i \in \{1, 2, \ldots, k{-}1\}$ it holds that $\bar{p}_i \cdot \bar{x}_{i+1} \leq \bar{p}_i \cdot \bar{x}_i$, then $\bar{p}_k \cdot \bar{x}_1 \geq \bar{p}_k \cdot \bar{x}_k$. It is straightforward from the definitions that satisfying GARP is a precondition for rationalizability (indeed, otherwise we would have for any rationalizing utility function u that $u(\bar{x}_1) \geq u(\bar{x}_2) \geq \cdots \geq u(\bar{x}_k) > u(\bar{x}_1)$). In a celebrated result, Afriat (1967) showed that GARP is also a sufficient condition for rationalizability of a finite dataset—and furthermore GARP is a sufficient condition for rationalizability of such a dataset by a utility function with many properties that are often assumed in simple economic models. This finding implies that the standard economic model of consumer choice has no testable implications beyond GARP.

Theorem 6.1 (Afriat, 1967)**.** *A finite dataset $S \subseteq \left(\mathbb{R}^m_+ \setminus \{\bar{0}\}\right) \times \mathbb{R}^m_+$ satisfies GARP if and only if it is rationalizable. Moreover, when GARP holds there exists a utility function that rationalizes S that is continuous, concave, nondecreasing, and strictly increases when all coordinates strictly increase.*

There are well-known examples of infinite datasets that are generated by quasiconcave utility functions but may not be rationalized by a concave utility function (see Aumann, 1975; Reny, 2013). Kannai (2004) and Apartsin and Kannai (2006) provide necessary conditions, stronger than GARP, for rationalizability by a concave function. Recently, Reny (2015) unified the literature and clarified the boundaries of Afriat's theorem when datasets are infinite by showing that GARP is indeed necessary and sufficient for rationalization of even infinite datasets—and in fact, GARP also guarantees rationalizability by a utility function with many desired properties (yet not all the properties that are attainable in the finite case).

Theorem 6.2 (Reny, 2015)**.** *A (possibly infinite) dataset $S \subseteq \left(\mathbb{R}^m_+ \setminus \{\bar{0}\}\right) \times \mathbb{R}^m_+$ satisfies GARP if and only if it is rationalizable. Moreover, when GARP holds there exists a utility function that rationalizes S that is quasiconcave and nondecreasing, and strictly increases when all coordinates strictly increase.*

Reny (2015) provides examples showing that continuity and concavity (the properties of the rationalizing utility function from Theorem 6.1 that are absent from Theorem 6.2) cannot be guaranteed to be attainable for any rationalizable dataset. Reny (2015) then proves Theorem 6.2 by deriving a novel way to construct a nondecreasing and quasiconcave utility function rationalizing a given finite data set that—unlike Afriat's construction—permits a lifting to infinite data sets.

[29] This assumption rules out trivial rationalizations such as constant utility functions. See Chambers and Echenique (2016) for a more detailed discussion.

We instead give a concise alternative proof[30] of Theorem 6.2 by lifting Theorem 6.1 as a black box using Logical Compactness. One of the challenges in our argument is that utility functions have an infinite range, so it is not *a priori* obvious how to encode such a function by a model defined via individually finite formulae; we overcome this challenge by having the model encode an entire sequence of discrete functions that converges to the utility function. In Section 4.3 we solved a similar obstacle by showing the existence of a discrete approximate solution using Logical Compactness, and then taking the exact solution to be the limit of a sequence of such approximate solutions, but this approach does not seem suitable for proving Theorem 6.2. Either of the limit approaches just described—the one in this section and the one in Section 4.3—in turn creates additional challenges, such as how to make sure, using only constraints on these discrete functions, that certain properties that are not maintained by limits (such as being strictly increasing) are satisfied by the limit function.

Proof of Theorem 6.2. Let $S \subseteq \left(\mathbb{R}_+^m \setminus \{\bar{0}\}\right) \times \mathbb{R}_+^m$ be a dataset that satisfies GARP. We will search for a utility function u with the required properties that rationalizes S whose range is $[0,1]$. For every $n \in \mathbb{N}$, we set $\varepsilon_n = 2^{-n}$, and for every x we denote by $\lfloor x \rfloor_{\varepsilon_n} = 2^{-n} \cdot \lfloor 2^n \cdot x \rfloor$ the rounding-down of x to the nearest multiple of ε_n. For every $n \in \mathbb{N}$, every $\bar{x} \in \mathbb{R}_+^m$, and every $v \in V_n = \{0, \varepsilon_n, 2 \cdot \varepsilon_n, \ldots, 1\}$, we introduce a variable $\texttt{utility}_{\bar{x},v}^n$ that will be TRUE in a model if and only if $\lfloor u(\bar{x}) \rfloor_{\varepsilon_n} = v$ for the utility function u that corresponds to the model. We define the set of formulae Φ as follows:

1. For every $n \in \mathbb{N}$ and every $\bar{x} \in \mathbb{R}_+^m$, we add the following (finite!) formula:

   $$\bigvee_{v \in V_n} \texttt{utility}_{\bar{x},v}^n,$$

   requiring that $\bar{x}$ have a rounded-down-to-ε_n utility.

2. For every $n \in \mathbb{N}$, every $\bar{x} \in \mathbb{R}_+^m$, and every distinct $v, w \in V_n$, we add the following formula:

   $$\texttt{utility}_{\bar{x},v}^n \to \neg\texttt{utility}_{\bar{x},w}^n,$$

   requiring that the rounded-down-to-ε_n utility from $\bar{x}$ be unique.

3. For every $n \in \mathbb{N}$, every $\bar{x} \in \mathbb{R}_+^m$, and every $v \in V_n$, we add the following formula:

   $$\texttt{utility}_{\bar{x},v}^n \to \left(\texttt{utility}_{\bar{x},v}^{n+1} \vee \texttt{utility}_{\bar{x},v+\varepsilon_{n+1}}^{n+1}\right),$$

   requiring that $\lfloor u(\bar{x}) \rfloor_{\varepsilon_n} = \lfloor \lfloor u(\bar{x}) \rfloor_{\varepsilon_{n+1}} \rfloor_{\varepsilon_n}$.

4. For every $n \in \mathbb{N}$, every $\bar{x}, \bar{y} \in \mathbb{R}_+^m$, every $\bar{z} \in \mathbb{R}_+^m$ that is a convex combination of $\bar{x}, \bar{y}$, and every $v, w \in V_n$, we add the following (finite) formula:

   $$\left(\texttt{utility}_{\bar{x},v}^n \wedge \texttt{utility}_{\bar{y},w}^n\right) \to \bigvee_{\substack{v' \in V_n: \\ v' \geq \min\{v,w\}}} \texttt{utility}_{\bar{z},v'}^n,$$

   requiring that the rounded-down-to-ε_n utility function be quasiconcave.

[30] While we prove the result of Reny (2015) in its full generality, it is worth noting that Reny's proof does not use the Axiom of Choice, while ours does to some extent. More specifically, the Compactness Theorem is equivalent (under ZF) to the Boolean Prime Ideal (BPI) theorem (equivalently, to the Ultrafilter Lemma), which is known to be a "weaker form of the Axiom of Choice" in the sense that ZF+BPI is strictly weaker than ZFC but strictly stronger than ZF (see, e.g., Halbeisen, 2017, Theorems 6.7 and 8.16).

5. For every $n \in \mathbb{N}$, every $\bar{x}, \bar{y} \in \mathbb{R}^m_+$ s.t. $\bar{x} \leq \bar{y}$, and every $v \in V_n$, we add the following (finite) formula:
$$\texttt{utility}^n_{\bar{x},v} \to \bigvee_{\substack{w \in V_n:\\ w \geq v}} \texttt{utility}^n_{\bar{y},w},$$
requiring that the rounded-down-to-ε_n utility function be nondecreasing.

6. Let $(\bar{q}^k_1, \bar{q}^k_2)^\infty_{k=1}$ be an enumeration of the countable set $\{(\bar{q}_1, \bar{q}_2) \in \mathbb{Q}^m \times \mathbb{Q}^m \mid \bar{q}_1 \ll \bar{q}_2\}$.[31] For every $k \in \mathbb{N}$ and for every $n > k$, we add the following (finite) formula:
$$\texttt{utility}^n_{\bar{q}^k_1,v} \to \bigvee_{\substack{w \in V_n:\\ w \geq v + 2^{-k-1}}} \texttt{utility}^n_{\bar{q}^k_2,w},$$
requiring that starting at some n the rounded-down-to-ε_n utility from $\bar{q}^k_2$ be greater by at least 2^{-k-1} than the rounded-down-to-ε_n utility from $\bar{q}^k_1$.

7. For every $n \in \mathbb{N}$, every given datapoint $(\bar{p}, \bar{x}) \in S$, every $\bar{y} \in \mathbb{R}^m_+$ s.t. $\bar{p} \cdot \bar{y} \leq \bar{p} \cdot \bar{x}$, and every $v \in V_n$, we add the following (finite) formula:
$$\texttt{utility}^n_{\bar{x},v} \to \bigvee_{\substack{w \in V_n:\\ w \leq v}} \texttt{utility}^n_{\bar{y},w},$$
requiring that the rounded-down-to-ε_n utility weakly rationalize S.

We first claim that every model that satisfies Φ corresponds to a utility function that rationalizes S and is quasiconcave, nondecreasing, and strictly increases when all coordinates strictly increase. Indeed, for every $\bar{x} \in \mathbb{R}^m_+$, for every $n \in \mathbb{N}$ let $v_n \in V_n$ be the value such that $\texttt{utility}^n_{\bar{x},v_n}$ is TRUE in the model (well defined by the first and second formula-types above), and define $u(\bar{x}) = \lim_{n\to\infty} v_n$ (well defined, e.g., by the third formula-type above since v_n is a Cauchy sequence). The resulting utility function u is a limit of nondecreasing quasiconcave functions (by the fourth and fifth formula-types above) that weakly rationalize the data (by the seventh formula-type above), and hence itself is a nondecreasing quasiconcave function that weakly rationalizes the data. Furthermore, for every $\bar{x}, \bar{y} \in \mathbb{R}^m_+$ s.t. $\bar{x} \ll \bar{y}$, there exist two rational number vectors "in between" them, i.e., there exists $k \in \mathbb{N}$ s.t. $\bar{x} \ll \bar{q}^k_1 \ll \bar{q}^k_2 \ll \bar{y}$. Therefore, we have that $u(\bar{x}) \leq u(\bar{q}^k_1) \leq u(\bar{q}^k_2) - 2^{-k-1} < u(\bar{q}^k_2) \leq u(\bar{y})$ (the second inequality stems from this inequality holding for almost all functions of which u is the limit, by the sixth formula-type above), so u strictly increases when all coordinates strictly increase. Finally, since u weakly rationalizes S and also strictly increases when all coordinates strictly increase, then u also rationalizes S. So, it is enough to show that Φ is satisfiable, and by the Compactness Theorem, it is enough to show that every finite subset $\Phi' \subseteq \Phi$ is satisfiable.

Let Φ' be a finite subset of Φ. We note in particular that there are only finitely many formulae of the above sixth and seventh types in Φ'. Since there are only finitely many formulae of the above seventh type, to satisfy Φ' we need to rationalize only a finite dataset—the datapoints corresponding to these finitely many seventh-type formulae. By Theorem 6.1, there exists a function u whose range is contained in $\mathbb{R}$ that rationalizes all these finitely many datapoints, that is concave (and in particular quasiconcave) and nondecreasing, and that strictly increases when all coordinates strictly increase. We start by defining $\bar{u}(\cdot) = 1/4 + \frac{\arctan(u(\cdot))}{2\pi}$. As this transformation is strictly monotone,

[31] For $\bar{x}, \bar{y} \in \mathbb{R}^m$, we write $\bar{x} \ll \bar{y}$ to denote that $\bar{x}$ is strictly less than $\bar{y}$ in each of the m coordinates.

the resulting function $\bar{u}$ still rationalizes the data, is quasiconcave and nondecreasing, and strictly increases when all coordinates strictly increase. Furthermore, the range of $\bar{u}$ is contained in $[0, 1/2]$. Now, for each of the finitely many pairs $(\bar{q}_1^k, \bar{q}_2^k)$ for which the corresponding sixth-type formula is in Φ', we "massage" $\bar{u}$ by creating a gap of size 2^{-k-1} in the range of $\bar{u}$ just below $\bar{u}(\bar{q}_2^k)$. Formally, we define:

$$\bar{u}_{\text{new}}(\bar{x}) = \begin{cases} \bar{u}_{\text{old}}(\bar{x}) & \bar{u}_{\text{old}}(\bar{x}) < \bar{u}_{\text{old}}(\bar{q}_2^k) \\ \bar{u}_{\text{old}}(\bar{x}) + 2^{-k-1} & \text{otherwise} \end{cases}$$

(note that this is once again a strictly monotone transformation, and note that $\bar{u}_{\text{new}}(\bar{q}_2^k) = \bar{u}_{\text{old}}(\bar{q}_2^k) + 2^{-k-1} > \bar{u}_{\text{old}}(\bar{q}_1^k) + 2^{-k-1} = \bar{u}_{\text{new}}(\bar{q}_1^k) + 2^{-k-1}$), before "massaging" it again (where $\bar{u}_{\text{old}}$ of the next step would be $\bar{u}_{\text{new}}$ of this step) to create a gap for the next pair $(\bar{q}_1^{k'}, \bar{q}_2^{k'})$ for which corresponding sixth-type formula is in Φ'. Since the sum of all the gaps that we create is less than one half, the range of the resulting function $\bar{u}$ after creating all gaps is contained in $[0, 1]$. To sum up, at the end of the iterative process, $\bar{u}$ is now a quasiconcave and nondecreasing function whose range is contained in $[0, 1]$ that rationalizes all datapoints for which the corresponding seventh-type formula is in Φ', and such that for every pair $(\bar{q}_1^k, \bar{q}_2^k)$ for which the corresponding sixth-type formula is in Φ', satisfies $\bar{u}(\bar{q}_2^k) \geq \bar{u}(\bar{q}_1^k) + 2^{-k-1}$. Using $\bar{u}$ we now construct a model for Φ' by setting, for each $\bar{x} \in \mathbb{R}_+^m$, for each $n \in \mathbb{N}$, and for each $v \in V_n$, the variable $\texttt{utility}_{\bar{x},v}^n$ to be TRUE if $v = \lfloor \bar{u}(x) \rfloor_{\varepsilon_n}$ and to be FALSE otherwise. It is straightforward to verify that this indeed is a model for Φ' (and, in fact, also for all first- to fifth-type formulae in Φ), so Φ' is satisfiable as required. □

Our concise proof of Theorem 6.2 relies heavily—due to the inherent need to make each formula finite—on the fact that quasiconcavity is maintained under weakly monotone transformations (such as various "condensing" and "gap" operations, and such as the operation of rounding-down to a grid). Our proof also sheds light on the need to relax the stronger, cardinal concavity condition to the weaker, ordinal quasiconcavity condition in order to generalize Afriat's result to infinite datasets. While our language allows us to require approximate versions of concavity on the sequence, if we add the requirement of monotonicity (formulae of sixth type), we cannot be assured that the utility will be bounded (hindering our ability to rely on type-one and type-two formulae). Here we can see again (as we have already seen, e.g., in Section 5), how certain conditions that may seem at first glance to be forced upon us merely by the use of Compactness (and more specifically, by the need to express the conditions for our solution via individually finite formulae) can in fact turn out to be essential for the result, in the sense that they cannot be dispensed with. Thus, it is possible that Compactness may also be useful for highlighting possible break points for generalizing theorems from finite to infinite domains.

6.2 Rationalizing Stochastic Demand

Fix a set X of alternatives. A *(stochastic choice) dataset* is a function

$$P : \left\{ (A, x) \in (2^X \setminus \{\emptyset\}) \times X \;\middle|\; A \in \mathcal{A} \;\;\&\;\; x \in A \right\} \to [0, 1]$$

such that $\mathcal{A} \subseteq \left\{ A \in 2^X \setminus \{\emptyset\} \;\middle|\; |A| < \infty \right\}$ and $\sum_{x \in A} P(A, x) = 1$ for every $A \in \mathcal{A}$. The dataset P is interpreted as probabilities with which different alternatives are chosen given various menus $\mathcal{A}$ from X. Probabilistic choice may emerge from random shocks to preferences over time, or represent fractions of deterministic choices in a population.[32]

[32] For a detailed discussion and a textbook treatment of this setting, see Chambers and Echenique (2016).

A dataset is *rationalizable* if there exists a probability measure ν over the space[33] of full orders over X such that $\Pr_\nu[x \succ A \setminus \{x\}] = P(A,x)$ for every (A,x) in the domain of P. A dataset P satisfies the *Axiom of Revealed Stochastic Preference* if for every n and every finite sequence (possibly with repetitions) $(A_i, x_i)_{i=1}^n$, where each (A_i, x_i) is in the domain of P,

$$\sum_{i=1}^{n} P(A_i, x_i) \leq \max_{\pi \in (\cup_{i=1}^{n} A_i)!} \sum_{i=1}^{n} \mathbf{1}_{[x_i \succ_\pi A_i \setminus \{x_i\}]},$$

where the factorial symbolizes the set of all possible permutations.

Theorem 6.3 (McFadden and Richter 1971, 1990)**.** *Let X be a finite set of items. A dataset P is rationalizable if and only if it satisfies the Axiom of Revealed Stochastic Preference.*

In this section, we will use Logical Compactness to lift Theorem 6.3 to prove the following.

Theorem 6.4. *Let X be a (possibly infinite) set of items. A dataset P is rationalizable if and only if it satisfies the Axiom of Revealed Stochastic Preference.*

This infinite setting and its economic importance have been discussed by Cohen (1980) and McFadden (2005). Cohen (1980) showed that the celebrated representation result of Falmagne (1978) using Block–Marschak polynomials (Block and Marschak, 1959) extends to infinite sets X if the definition of rationalizability is weakened; Cohen (1980) also gave several stronger structural conditions on X that are sufficient for ("un-weakened," i.e., as defined above) rationalizability. McFadden (2005) showed how to extend Theorem 6.3 to an infinite setting different than ours (and different from the setting we study, which is the setting of Cohen, 1980), once again by either weakening the definition of rationalizability or by demanding the existence of a certain topological structure on X (to obtain "un-weakened" rationalizability). Theorem 6.4, which we will now prove, does *not* weaken the definition of rationalizability and does *not* impose any assumptions whatsoever on X.

Proof of Theorem 6.4. In our proof, we will use the following lemma to help us encode via individually finite formulae a probability measure μ over the full orders of X. The lemma, whose proof we spell out in Appendix A, follows directly from the Kolmogorov Extension Theorem.

Lemma 6.5. *Let X be a (possibly infinite) set, and for every $n \in \mathbb{N}$ and sequence $\bar{a} = (a_i)_{i=1}^n$ of n distinct elements of X, let $p_{\bar{a}} \in [0,1]$. Then the following are equivalent:*

- *There exists a probability measure μ over the space of full orders over X such that for every $n \in \mathbb{N}$ and sequence $\bar{a} = (a_i)_{i=1}^n$ of n distinct elements of X, it is the case that*

$$p_{\bar{a}} = \Pr_\mu[a_1 \succ a_2 \succ \cdots \succ a_n].$$

- $p_{(a)} = 1$ *for every $a \in X$ (sequence of length 1), and for every $n \in \mathbb{N}$ and sequence of $n+1$ distinct elements of X, $(a_1, \ldots, a_n, a)$ it is the case that* $p_{(a_1,\ldots,a_n)} = \sum_{i=0}^{n} p_{(a_1,\ldots,a_i,a,a_{i+1},\ldots,a_n)}$.

[33] In this section, when we consider probability measures over the space of full orders over a set X, then if X is countable then we will take this space as a measurable space w.r.t. the discrete σ-algebra, and more generally for arbitrary X we will take this space as a measurable space w.r.t. the σ-algebra generated by all of its subset of the form $\{\pi \mid a_1 \succ_\pi \cdots \succ_\pi a_n\}$, where the a_i are distinct elements in X.

Now, it is clear (the proof is immediate, as in the finite case), that satisfying the Axiom of Revealed Stochastic Preference is a necessary condition for rationalizability of P. We will prove that this condition also suffices for rationalizability. That is, given a dataset P that satisfies the Axiom of Revealed Stochastic Preference, we will prove the existence of a probability measure μ on the space of full orders over X that rationalizes P.

For every $n \in \mathbb{N}$, we set $\varepsilon_n = 2^{-n}$, and for every x we denote by $\lfloor x \rfloor_{\varepsilon_n} = 2^{-n} \cdot \lfloor 2^n \cdot x \rfloor$ the rounding-down of x to the nearest multiple of ε_n. For every $n \in \mathbb{N}$, every $m \in \mathbb{N}$, every m-tuple of distinct items $\bar{a} = (a_1, \ldots, a_n)$ from X, and every $p \in V_n = \{0, \varepsilon_n, 2 \cdot \varepsilon_n, \ldots, 1\}$, we will have a variable $\mathtt{prob}^n_{\bar{a},p}$ that will be TRUE in a model if and only if $p = \lfloor \Pr_\mu[a_1 \succ a_2 \succ \cdots \succ a_m] \rfloor_{\varepsilon_n}$ for the probability measure μ that corresponds to the model. We will define the set of formulae Φ as follows:

1. For every $n \in \mathbb{N}$ and every tuple $\bar{a}$ of distinct items from X, we add the following (finite!) formula:
$$\bigvee_{v \in V_n} \mathtt{prob}^n_{\bar{a},p},$$
requiring that $\bar{a}$ have a rounded-down-to-ε_n probability.

2. For every $n \in \mathbb{N}$, every tuple $\bar{a}$ of distinct items from X, and every distinct $p, q \in V_n$, we add the following formula:
$$\mathtt{prob}^n_{\bar{a},p} \to \neg \mathtt{prob}^n_{\bar{a},q},$$
requiring that the rounded-down-to-ε_n probability of $\bar{a}$ be unique.

3. For every $n \in \mathbb{N}$, every tuple $\bar{a}$ of distinct items from X, and every $p \in V_n$, we add the following formula:
$$\mathtt{prob}^n_{\bar{a},p} \to \left(\mathtt{prob}^{n+1}_{\bar{a},p} \vee \mathtt{prob}^{n+1}_{\bar{a},p+\varepsilon_{n+1}} \right),$$
requiring that $\lfloor \Pr_\mu[a_1 \succ a_2 \succ \cdots \succ a_m] \rfloor_{\varepsilon_n} = \lfloor \lfloor \Pr_\mu[a_1 \succ a_2 \succ \cdots \succ a_m] \rfloor_{\varepsilon_{n+1}} \rfloor_{\varepsilon_n}$.

4. For every $n \in \mathbb{N}$ and every $a \in X$, we add the following (finite) formula:
$$\mathtt{prob}^n_{(a),1},$$
requiring that the rounded-down-to-ε_n probability of the ordering a is 1.

5. For every $n \in \mathbb{N}$, every $m \in \mathbb{N}$, and every $(m+1)$-tuple $(a_1, \ldots, a_m, a)$ of distinct items from X, we add the following (finite) formula:
$$\bigvee_{\substack{p, p_0, \ldots, p_m \in V_n \\ \text{s.t. } \sum_{i=1}^{m+1} p_i \in [p - (m+1) \cdot \varepsilon_n, p]}} \left(\mathtt{prob}^n_{(a_1, \ldots, a_m), p} \wedge \bigwedge_{i=0}^{m} \mathtt{prob}^n_{(a_1, \ldots, a_i, a, a_{i+1}, \ldots, a_m), p_i} \right),$$
requiring that up to rounding errors, the second condition of Lemma 6.5 hold for every "rounding-down of μ."

6. For every $n \in \mathbb{N}$ and every (A, x) in the domain of P, we add the following (finite) formula:
$$\bigvee_{\substack{p_1, \ldots, p_{(|A|-1)!} \in V_n \\ \text{s.t. } \sum_{i=1}^{(|A|-1)!} p_i \in \\ [P(A,x) - (|A|-1)! \cdot \varepsilon_n \,,\, P(A,x)]}} \bigwedge_{i=1}^{(|A|-1)!} \mathtt{prob}^n_{\underbrace{(x, a_1, \ldots, a_{|A|-1})}_{\substack{i\text{th permutation} \\ \text{of } A \setminus \{x\}}}, p_i},$$
requiring that up to rounding errors, every "rounding-down of μ" rationalize P.

We first claim that every model that satisfies Φ corresponds to a probability measure that rationalizes P. Indeed, for every $m \in \mathbb{N}$ and m-tuple $\bar{a}$ of distinct items from X, for every $n \in \mathbb{N}$ let $p_n \in V_n$ be the probability such that $\mathtt{prob}^n_{\bar{a},p_n}$ is TRUE in the model (well defined by the first and second formula-types above), and define $p_{\bar{a}} = \lim_{n\to\infty} p_n$ (well defined, e.g., by the third formula-type above since $\mathtt{prob}^n_{\bar{a},p_n}$ is a Cauchy sequence). The resulting probabilities $p_{\bar{a}}$ satisfy the second condition of Lemma 6.5 (by the fourth and fifth formula-types above), and hence there exists a probability measure μ over the space of full orders over X that induces these probabilities. Finally, since μ induces these probabilities, then by the sixth formula-type above, μ rationalizes P. So, it is enough to show that Φ is satisfiable, and by the Compactness Theorem, it is enough to show that every finite subset $\Phi' \subseteq \Phi$ is satisfiable.

Each finite subset Φ' of Φ mentions only a finite set of items $X' \subset X$. Let P' be the restriction of P to all (A, x) in its domain such that $A \subseteq X'$. By Theorem 6.3, there exists a probability measure μ over $X'!$ that rationalizes P'. Using μ we now construct a model for Φ' by setting, for each $m \in \{1, \ldots, |X'|\}$ and m-tuple $\bar{a}$ of distinct items from X', for each $n \in \mathbb{N}$, and for each $p \in V_n$, the variable $\mathtt{prob}^n_{\bar{a},p}$ to be TRUE if $p = \lfloor \Pr_\mu[a_1 \succ a_2 \succ \cdots \succ a_m] \rfloor_{\varepsilon_n}$ and to be FALSE otherwise. It is straightforward to verify that this indeed is a model for Φ' (and, in fact, also for all other formulae from Φ that only mention items in X'), so Φ' is satisfiable as required. $\square$

7 Conclusion

Propositional Logic gives a principled way to extend economic theory results from finite models to infinite ones. The resulting arguments are intuitive and (of course) compact.

As we have demonstrated, the Compactness Theorem for Propositional Logic can be used to scale economic theory models from finite to infinite in a range of different ways. We first showed how Logical Compactness lets us extend results from finite matching markets and finite trading networks to large-market settings; we think of the results that extend in this way as being especially robust, in the sense that they do not rely on edge effects or specific starting conditions. For example, while the Lone Wolf/Rural Hospitals Theorem in matching markets is usually considered more fundamental than strategy-proofness (and indeed the former is used in the standard proof of the latter), we show that the latter holds in infinite markets even though the former is known to break in such markets due to edge effects, showing that in some sense the latter is in fact more robust. Moreover, in the large-market equilibria and stable matchings that we obtain using Logical Compactness, agents "maintain their mass," in the sense that they are still subject to strategic incentives, and even individually affect others' strategic incentives; thus we might think of these results as providing especially realistic large-market models of economic behavior.

At the same time, we showed that Logical Compactness also lets us extend repeated economic interactions with finite start-times into "ongoing" dynamic interactions with neither start nor end—providing a model of a steady-state free from behavioral artefacts driven by time starting or ending at a fixed date, in a formal sense allowing us to perform induction without a base for the induction, as we have done for dynamic matching.

Last, we took Logical Compactness beyond existence and characterization of solutions, to extend decision-theoretic results from finite-data settings to infinite-data ones. This allows us to reason about how an agent's choice would behave on arbitrary input data. For example, let P be a finite set of price vectors and $P' \supset P$ be a larger set of price vectors; and say that we are given the demand of a rational consumer at each price vector in P, and are interested in the possible demands not only over P but also over P'. If P' is finite, then Afriat's Theorem tells us that any demand over P' that is consistent with the given demand over P and satisfied GARP is possible. The infinite-data version of Afriat's Theorem tells us the same even if P' is infinite. So, even given only finite data, an infinite-data version of this theorem was needed to be able to reason about an agent's future choice. Moreover, such infinite-data extensions also give us a clearer understanding of fundamental limitations on inference about agents' preferences. For example, showing that the inability of demand queries to rule out continuity and concavity of the utility function, as indicated by Afriat's Theorem, is an artefact of finite data—as it is possible to rule these out given infinite data—while the inability of demand queries to rule out quasiconcavity of the utility function persists even with infinite data, and is thus a fundamental limitation.

While the Compactness Theorem is implied by more general results in topology, we also used our approach to give intuitive proofs of results (such as strategy-proofness) that do not seem inherently topological in nature. From a methodological perspective, Logical Compactness has the advantage that it (1) reduces reasoning about infinite problems to reasoning about finite problems, and then (2) allows us to think about those finite problems in their natural language. We hope this will make Logical Compactness both easily and flexibly applicable in a range of theory contexts even beyond those we have considered here.

References

S. N. Afriat. The construction of utility functions from expenditure data. *International Economic Review*, 8(1):67–77, 1967.

V. Aguiar, P. Hjertstrand, and R. Serrano. A rationalization of the weak axiom of revealed preference. Working paper, Bravo Center for Economic Research, 2020.

S. N. Ali and C. Liu. Conventions and coalitions in repeated games. Mimeo, 2019.

R. M. Anderson. An elementary core equivalence theorem. *Econometrica: Journal of the Econometric Society*, pages 1483–1487, 1978.

R. M. Anderson. Non-standard analysis with applications to economics. In W. Hildenbrand and H. Sonnenschein, editors, *Handbook of Mathematical Economics*, volume 4, chapter 39, pages 2145–2208. Elsevier, 1991.

Y. Apartsin and Y. Kannai. Demand properties of concavifiable preferences. *Journal of Mathematical Economics*, 43(1):36–55, 2006.

I. Arieli and R. J. Aumann. The logic of backward induction. *Journal of Economic Theory*, 159: 443–464, 2015.

I. Ashlagi, M. Braverman, and A. Hassidim. Stability in large matching markets with complementarities. *Operations Research*, 62(4):713–732, 2014.

R. J. Aumann. Values of markets with a continuum of traders. *Econometrica: Journal of the Econometric Society*, pages 611–646, 1975.

R. J. Aumann and L. S. Shapley. *Values of Non-Atomic games*. Princeton University Press, 1974.

E. M. Azevedo and E. Budish. Strategy-proofness in the Large. *Review of Economic Studies*, 86 (1):81–116, 2018.

E. M. Azevedo and J. W. Hatfield. Existence of stable matchings in large markets with complementarities. Working paper, University of Texas at Austin, 2015.

E. M. Azevedo and J. D. Leshno. A supply and demand framework for two-sided matching markets. *Journal of Political Economy*, 124:1235–1268, 2016.

E. M. Azevedo, E. G. Weyl, and A. White. Walrasian equilibrium in large, quasilinear markets. *Theoretical Economics*, 8:281–290, 2013.

A. Bjorndahl, J. Y. Halpern, and R. Pass. Language-based games. In *Proceedings of 14th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 39–48, 2014.

H. D. Block and J. Marschak. Random orderings and stochastic theories of response. Technical report, Cowles Foundation for Research in Economics, 1959.

L. E. Blume and W. R. Zame. The algebraic geometry of perfect and sequential equilibrium. *Econometrica: Journal of the Econometric Society*, 62(4):783–794, 1994.

D. J. Brown and M. A. Khan. An extension of the Brown-Robinson equivalence theorem. *Applied Mathematics and Computation*, 6(2):167–175, 1980.

C. P. Chambers and F. Echenique. *Revealed Preference Theory*. Cambridge University Press, 2016.

C. P. Chambers, F. Echenique, and E. Shmaya. The axiomatic structure of empirical content. *American Economic Review*, 104(8):2303–19, 2014.

C. P. Chambers, F. Echenique, and E. Shmaya. General revealed preference theory. *Theoretical Economics*, 12(2):493–511, 2017.

Y.-K. Che, J. Kim, and F. Kojima. Stable matching in large economies. *Econometrica: Journal of the Econometric Society*, 87(1):65–110, 2019.

Y. Choi. On two-sided matching in infinite markets. Working paper, Research Science Institute, 2020.

D. Clark, D. Fudenberg, and A. Wolitzky. Robust cooperation with first-order information. Mimeo, 2019.

M. A. Cohen. Random utility systems—the infinite case. *Journal of Mathematical Psychology*, 22 (1):1–23, 1980.

N. G. de Bruijn and P. Erdős. A colour problem for infinite graphs and a problem in the theory of relations. *Indigationes Mathematicae*, 13(5):371–373, 1951.

L. Doval. A theory of stability in dynamic matching markets. Mimeo, 2017.

L. E. Dubins and D. A. Freedman. Machiavelli and the Gale-Shapley algorithm. *American Mathematical Monthly*, 88:485–494, 1981.

J. Duggan. A general extension theorem for binary relations. *Journal of Economic Theory*, 86(1): 1–16, 1999.

J.-C. Falmagne. A representation theorem for finite random scale systems. *Journal of Mathematical Psychology*, 18(1):52–72, 1978.

Y. Feinberg. Characterizing common priors in the form of posteriors. *Journal of Economic Theory*, 91:127–179, 2000.

P. C. Fishburn. Comment on the Kannai-Peleg impossibility theorem for extending orders. *Journal of Economic Theory*, 32(1):176–179, 1984.

T. Fleiner. A fixed-point approach to stable matchings and some applications. *Mathematics of Operations Research*, 28:103–126, 2003.

D. Fudenberg and D. K. Levine. *A Long-Run Collaboration on Long-Run Games*. World Scientific, 2009.

M. Fuentes and F. Tohmé. Stable matching with double infinity of workers and firms. *The BE Journal of Theoretical Economics*, 19(2), 2019.

D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15, 1962.

M. Greinecker and C. Kah. Pairwise stable matching in large economies. Mimeo, 2019.

N. E. Gretsky, J. M. Ostroy, and W. R. Zame. The nonatomic assignment model. *Economic Theory*, 2(1):103–127, 1992.

N. E. Gretsky, J. M. Ostroy, and W. R. Zame. Perfect competition in the continuous assignment model. *Journal of Economic Theory*, 88(1):60–118, 1999.

L. J. Halbeisen. *Combinatorial Set Theory With a Gentle Introduction to Forcing*. Springer Monographs in Mathematics. Springer, second edition, 2017.

P. Hall. On representatives of subsets. *Journal of the London Mathematical Society*, 1(1):26–30, 1935.

P. R. Halmos and H. E. Vaughan. The marriage problem. *American Journal of Mathematics*, 72 (1):214–215, 1950.

J. Y. Halpern. A nonstandard characterization of sequential equilibrium, perfect equilibrium, and proper equilibrium. *International Journal of Game Theory*, 38:37–49, 2009. Erratum: *International Journal of Game Theory*, 46:591–594, 2017.

J. Y. Halpern. Lexicographic probability, conditional probability, and nonstandard probability. *Games and Economic Behavior*, 68(1):155–179, 2010.

J. Y. Halpern and Y. Moses. Characterizing solution concepts in terms of common knowledge of rationality. *International Journal of Game Theory*, 46:457–473, 2016.

J. W. Hatfield and S. D. Kominers. Hidden substitutes. Harvard University Working Paper, 2019.

J. W. Hatfield, S. D. Kominers, A. Nichifor, M. Ostrovsky, and A. Westkamp. Stability and competitive equilibrium in trading networks. *Journal of Political Economy*, 121(5):966–1005, 2013.

J. W. Hatfield, S. D. Kominers, A. Nichifor, M. Ostrovsky, and A. Westkamp. Full substitutability. *Theoretical Economics*, 14(4):1535–1590, 2019.

Z. Hellman and Y. J. Levy. Measurable selection for purely atomic games. *Econometrica: Journal of the Econometric Society*, 87(2):593–629, 2019.

R. Holzman. An extension of fishburn's theorem on extending orders. *Journal of Economic Theory*, 32(1):192–196, 1984.

N. Immorlica and M. Mahdian. Marriage, honesty, and stability. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 53–62, 2005.

R. Jagadeesan. Complementary inputs and the existence of stable outcomes in large trading networks. Working paper, Harvard University, 2018a.

R. Jagadeesan. Lone wolves in infinite, discrete matching markets. *Games and Economic Behavior*, 108:275–286, 2018b.

S. V. Kadam and M. H. Kotowski. Multiperiod matching. *International Economic Review*, 59(4): 1927–1947, 2018.

Y. Kamada and F. Kojima. Efficient matching under distributional constraints: Theory and applications. *American Economic Review*, 105:67–99, 2015.

M. Kaneko and M. H. Wooders. The core of a game with a continuum of players and finite coalitions: The model and some results. *Mathematical Social Sciences*, 12(2):105–137, 1986.

Y. Kannai. When is individual demand concavifiable? *Journal of Mathematical Economics*, 40 (1-2):59–69, 2004.

M. Kearns. Graphical games. In N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 9, pages 159–178. Cambridge University Press, 2007.

M. A. Khan. The irony in/of economic theory. *MLN*, 108(4):759–803, 1993.

P. Kircher. Efficiency of simultaneous search. *Journal of Political Economy*, 117(5):861–913, 2009.

B. Klaus and F. Klijn. Stable matchings and preferences of couples. *Journal of Economic Theory*, 121:75–106, 2005.

F. Kojima and P. A. Pathak. Incentives and stability in large two-sided matching markets. *American Economic Review*, 99:608–627, 2009.

F. Kojima, P. A. Pathak, and A. E. Roth. Matching with couples: Stability and incentives in large markets. *Quarterly Journal of Economics*, 128(4):1585–1632, 2013.

S. Lahiri. A simple proof of Suzumura's extension theorem for finite domains with applications. *Journal of Applied Mathematics and Decision Sciences*, 6(3):183–190, 2002.

C. Liu. Stability in repeated matching markets. Mimeo, 2018.

D. Marker. *Model theory: an introduction*, volume 217. Springer Science & Business Media, 2006.

A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*. Oxford University Press, 1995.

M. Maschler, E. Solan, and S. Zamir. *Game Theory*. Cambridge University Press, 2013.

D. McFadden and M. Richter. On the extension of a set function on a set of events to a probability on the generated boolean σ-algebra. Working paper, University of California, Berkeley, 1971.

D. McFadden and M. K. Richter. Stochastic rationality and revealed stochastic preference. In J. Chipman, D. McFadden, and M. K. Richter, editors, *Preferences, Uncertainty, and Optimality, Essays in Honor of Leo Hurwicz*, pages 161–186. Westview Press, 1990.

D. L. McFadden. Revealed stochastic preference: a synthesis. *Economic Theory*, 26(2):245–264, 2005.

A. Miralles and M. Pycia. Large vs. continuum assignment economies. Mimeo, 2015.

J. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.

T. Nguyen and R. Vohra. Near-feasible stable matchings with couples. *American Economic Review*, 108(11):3154–69, 2018.

G. Nöldeke and L. Samuelson. The implementation duality. *Econometrica: Journal of the Econometric Society*, 86(4):1283–1324, 2018.

E. A. Ok. *Real Analysis with Economic Applications*, volume 10. Princeton University Press, 2007.

A. Öry. Consumers on a leash: Advertised sales and intertemporal price discrimination. Discussion paper, Cowles Foundation, 2016.

B. Peleg. Equilibrium points for games with infinitely many players. *Journal of the London Mathematical Society*, 44(1):292–294, 1969.

J. S. Pereyra. A dynamic school choice model. *Games and Economic Behavior*, 80:100–114, 2013.

P. J. Reny. A simple proof of the nonconcavifiability of functions with linear not-all-parallel contour sets. *Journal of Mathematical Economics*, 49(6):506–508, 2013.

P. J. Reny. A characterization of rationalizable consumer behavior. *Econometrica: Journal of the Econometric Society*, 83(1):175–192, 2015.

M. K. Richter. Revealed preference theory. *Econometrica: Journal of the Econometric Society*, 34 (3):635, 1966.

A. E. Roth. The economics of matching: Stability and incentives. *Mathematics of Operations Research*, 7:617–628, 1982.

A. E. Roth. The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy*, 92(6):991–1016, 1984.

A. E. Roth and M. Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, volume 18 of *Econometric Society Monographs*. Cambridge University Press, 1990.

H. E. Scarf. The core of an N person game. *Econometrica: Journal of the Econometric Society*, 35 (1):50–69, 1967.

L. Shapley and H. Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1(1): 23–37, 1974.

E. Szpilrajn. Sur l'extension de l'ordre partiel. *Fundamenta Mathematicae*, 1(16):386–389, 1930.

A Proof of Lemma 6.5

Proof of Lemma 6.5. The first condition immediately implies the second condition. We will therefore assume the second condition and prove the first condition from it. Let $T = \binom{X}{2}$, where we will denote an element of T as (i_1, j_1) with the choice of the ordering of i_1 and j_1 consistent throughout this proof for each pair. For every distinct $(i_1, j_1), \ldots, (i_k, j_k) \in T$, let $\nu_{(i_1,j_1),\ldots,(i_k,j_k)}$ be the probability measure over $\mathbb{R}^k$ defined as follows: Let $E = \{i_1, j_1, i_2, j_2, \ldots, i_k, j_k\}$ and let $n = |E| \leq 2k$. We define a probability measure μ_E over the finite set $E!$ of all permutations of E by assigning probability $p_{(e_1,\ldots,e_n)}$ to the ordering $(e_1, \ldots, e_n)$. By the second condition of the lemma, this is indeed a probability measure. We define $\nu_{(i_1,j_1),\ldots,(i_k,j_k)}$ as follows: to draw $r_1, \ldots, r_k \in \mathbb{R}^k$ according to $\nu_{(i_1,j_1),\ldots,(i_k,j_k)}$, first draw a permutation in $E!$ according to μ_E, and then for every $\ell \in \{1, \ldots, k\}$, set $r_\ell = 1$ if i_ℓ precedes j_ℓ according to this permutation, and $r_\ell = 0$ otherwise. Notice that for every measurable $F_1, \ldots, F_k \subseteq \mathbb{R}$, both of the following hold:

- For every permutation π of $\{1, \ldots, k\}$, we have that $\nu_{\pi((i_1,j_1)),\ldots,\pi((i_k,j_k))}(F_{\pi(1)}, \ldots, F_{\pi(k)}) = \nu_{(i_1,j_1),\ldots,(i_k,j_k)}(F_1, \ldots, F_k)$, immediately by the definition of these probability measures.
- For every $m \in \mathbb{N}$ and $(i_{k+1}, j_{k+1}), \ldots, (i_{k+m}, j_{k+m}) \in T$ distinct from one another and from $(i_1, j_1), \ldots, (i_k, j_k)$, it is the case that $\nu_{(i_1,j_1),\ldots,(i_m,j_m)}(F_1, \ldots, F_k, \mathbb{R}, \ldots, \mathbb{R}) = \nu_{(i_1,j_1),\ldots,\ldots,(i_k,j_k)}(F_1, \ldots, F_k)$, by the second condition of the lemma.

By these two conditions and by the Kolmogorov Extension Theorem, there exists a probability measure ν over $\mathbb{R}^T$ with the product σ-algebra whose marginals are the above-defined $\nu_{(\cdot)}$ measures.

To define the required probability measure μ, consider the following embedding into $\mathbb{R}^T$ of the space of full orders π over X: map a full order π over X to $(r_t)_{t \in T}$, where for every $(i, j) \in T$ we set $r_{(i,j)} = 1$ if i precedes j according to π, and $r_{(i,j)} = 0$ otherwise. We note that this embedding is an isomorphism of measurable spaces (i.e., a measurable bijection whose inverse is also measurable) of its domain (w.r.t. the σ-algebra generated by all of its subsets of the form $\{\pi \mid a_1 \succ_\pi \cdots \succ_\pi a_n\}$ where the a_is are distinct elements in X) and its image (w.r.t. the product σ-algebra), and therefore via this embedding the measure ν induces a measure μ over the space of full orders over (X w.r.t. to the above-defined σ-algebra). We note that the complement of the image of this embedding, inside $\mathbb{R}^T$, has measure 0 w.r.t. ν by the above construction of the marginals of ν, and so μ is a probability measure. By the definition of the marginals of ν via the μ_E measures defined above, the probability measure μ satisfies the first condition of the lemma, as required. $\square$

B Stable Matchings with Couples

Proof of Theorem 4.7. Assume without loss of generality that the preference lists of couples are such that no couple ranks any $(h, h') \in H \times H$ below $(h, \emptyset)$ or $(\emptyset, h')$. For every doctor $d \in D$ (single or from a couple) and hospital $h \in H$ we will have a variable $\texttt{matched}_{(d,h)}$ that will be TRUE in a model if and only if d and h are matched (in the matching corresponding to the model). For convenience, for every doctor d from a couple we will also have a variable $\texttt{matched}_{(d,\emptyset)}$ that will be TRUE in a model if and only if d is unmatched while d's partner is matched (see below). Furthermore, for every hospital $h \in H$ with capacity k_h we will have five variables $\texttt{capacity}_{(h,k_h-2)}, \texttt{capacity}_{(h,k_h-1)}, \ldots, \texttt{capacity}_{(h,k_h+2)}$ such that $\texttt{capacity}_{(h,q)}$ will be TRUE in a model if and only if $k^* = q$. We now proceed to define the set of formulae:

- Allowed adjusted capacities:

1. For every hospital h we add the following formula:

$$\bigvee_{q=k_h-2}^{k_h+2} \texttt{capacity}_{(h,q)},$$

requiring that h have an adjusted capacity of $k_h \pm 2$.

2. For every hospital h and two adjusted capacities $q \neq q'$ in $\{k_h - 2, \ldots, k_h + 2\}$ we add the following formula:

$$\texttt{capacity}_{(h,q)} \to \neg\texttt{capacity}_{(h,q')},$$

requiring that h have no more than one adjusted capacity.

- Doctor-to-hospital matching respecting adjusted capacities:

3. For every doctor d and for every two hospitals $h \neq h'$, we add the following formula:

$$\texttt{matched}_{(d,h)} \to \neg\texttt{matched}_{(d,h')},$$

requiring that d be matched to at most one hospital.

4. For every hospital h, possible capacity $q \in \{k_h - 2, \ldots, k_h + 2\}$, and $q+1$ distinct doctors $d, d_1, d_2, \ldots, d_q$, we add the following formula:

$$\left(\texttt{capacity}_{(h,q)} \wedge \bigwedge_{i=1}^{q} \texttt{matched}_{(d_i,h)}\right) \to \neg\texttt{matched}_{(d,h)},$$

requiring that h is not matched to more doctors than its capacity.

- Individual rationality:

5. For every doctor d and hospital h such that one or more of the following holds:
   (a) h does not rank d.
   (b) d is in D^1 and does not rank h.
   (c) d is in a couple, and no pair of hospitals in this couple's preference list has h matched to d.

   we add the following formula:

$$\neg\texttt{matched}_{(d,h)},$$

requiring that no doctor or hospital is matched in a way that they individually find unacceptable.

6. For every couple c and hospitals h, h' such that c that does not rank (h, h'), we add the following formula:

$$\neg(\texttt{matched}_{(c_1,h)} \wedge \texttt{matched}_{(c_2,h)}),$$

requiring that c are not matched to a pair of hospitals that they find unacceptable.

7a. For every couple c and hospital h such that c ranks $(h, \emptyset)$, let $h_1, \ldots, h_n$ be the hospitals such that c ranks every (h, h_i). We add the following formula:

$$\texttt{matched}_{(c_1,h)} \to \left(\texttt{matched}_{(c_2,\emptyset)} \leftrightarrow \neg \bigvee_{i=1}^{n} \texttt{matched}_{(c_2,h_i)}\right),$$

effectively setting (for convenience),when c_1 is matched with h, $\texttt{matched}_{(c_2,\emptyset)}$ as a shorthand for c_2 not being matched to any of these h_is.

7b. Completely symmetrically, for every couple c and hospital h such that c ranks $(\emptyset, h)$, let $h_1, \dots, h_n$ be the hospitals such that c ranks every (h_i, h). We add the following formula:

$$\texttt{matched}_{(c_2,h)} \to \left(\texttt{matched}_{(c_1,\emptyset)} \leftrightarrow \neg \bigvee_{i=1}^{n} \texttt{matched}_{(c_1,h_i)}\right),$$

effectively setting (for convenience), when c_2 is matched with h, $\texttt{matched}_{(c_1,\emptyset)}$ as a shorthand for c_1 not being matched to any of these h_is.

- Not blocked with respect to adjusted capacities:

8. For every single doctor $d \in D^1$ and hospital h that mutually rank each other, and for every possible capacity $q \in \{k_h - 2, \dots, k_h + 2\}$ for h, let $h_1, \dots, h_l$ be all the hospitals that d prefers to h and let $D_1, \dots, D_n$ be all q-tuples of doctors that h prefers to d. We add the following (finite!) formula:

$$(\texttt{capacity}_{(h,q)} \wedge \neg\texttt{matched}_{(d,h)}) \to \left(\bigvee_{i=1}^{l} \texttt{matched}_{(d,h_i)} \vee \bigvee_{i=1}^{n} \bigwedge_{d' \in D_i} \texttt{matched}_{(d',h)}\right),$$

which requires that (d, h) is not a blocking pair.

9. For every couple $c \in D^2$ and every two (actual) hospitals $h \neq h'$ such that c ranks (h, h') and such that h ranks c_1 and h' ranks c_2, and for every possible capacities $q \in \{k_h - 2, \dots, k_h+2\}$ and $q' \in \{k_{h'}-2, \dots, k_{h'}+2\}$ for h and h' respectively, let $(h_1, h'_1) \dots, (h_l, h'_l)$ be all the assignments that c prefers to (h, h') and let $D_1, \dots, D_n$ be all q-tuples of doctors that h prefers to c_1 and $D'_1, \dots, D'_{n'}$ be all q'-tuples of doctors that h' prefers to c_2. We add the following (finite) formula:

$$\begin{aligned}&(\texttt{capacity}_{(h,q)} \wedge \texttt{capacity}_{(h',q')} \wedge \neg(\texttt{matched}_{(c_1,h)} \wedge \texttt{matched}_{(c_2,h')})) \to \\ &\to \left(\bigvee_{i=1}^{l} (\texttt{matched}_{(c_1,h_i)} \wedge \texttt{matched}_{(c_2,h'_i)}) \vee \bigvee_{i=1}^{n} \bigwedge_{d \in D_i} \texttt{matched}_{(d,h)} \vee \bigvee_{i=1}^{n'} \bigwedge_{d' \in D'_i} \texttt{matched}_{(d',h')}\right),\end{aligned}$$

which requires that c is not blocking with (h, h').

10a. For every couple $c \in D^2$ and every hospital h such that c ranks $(h, \emptyset)$ and such that h ranks c_1, and for every possible capacity $q \in \{k_h - 2, \dots, k_h + 2\}$ for h, let $(h_1, h'_1) \dots, (h_l, h'_l)$ be all the assignments that c prefers to $(h, \emptyset)$ and let $D_1, \dots, D_n$ be all q-tuples of doctors that h prefers to c_1. We add the following (finite) formula:

$$\begin{aligned}&(\texttt{capacity}_{(h,q)} \wedge \neg(\texttt{matched}_{(c_1,h)} \wedge \texttt{matched}_{(c_2,\emptyset)})) \to \\ &\qquad \to \left(\bigvee_{i=1}^{l} (\texttt{matched}_{(c_1,h_i)} \wedge \texttt{matched}_{(c_2,h'_i)}) \vee \bigvee_{i=1}^{n} \bigwedge_{d \in D_i} \texttt{matched}_{(d,h)}\right),\end{aligned}$$

which requires that c is not blocking with $(h, \emptyset)$.

10b. Completely symmetrically, for every couple $c \in D^2$ and every hospital h such that c ranks $(\emptyset, h)$ and such that h ranks c_2, and for every possible capacity $q \in \{k_h-2, \dots, k_h+2\}$ for

h, let $(h_1, h'_1) \ldots, (h_l, h'_l)$ be all the assignments that c prefers to $(\emptyset, h)$ and let $D_1, \ldots, D_n$ be all q-tuples of doctors that h prefers to c_2. We add the following (finite) formula:

$$(\texttt{capacity}_{(h,q)} \wedge \neg(\texttt{matched}_{(c_1,\emptyset)} \wedge \texttt{matched}_{(c_2,h)})) \to$$
$$\to \left(\bigvee_{i=1}^{l} (\texttt{matched}_{(c_1,h_i)} \wedge \texttt{matched}_{(c_2,h'_i)}) \vee \bigvee_{i=1}^{n} \bigwedge_{d \in D_i} \texttt{matched}_{(d,h)} \right),$$

which requires that c is not blocking with $(\emptyset, h)$.

11. For every couple $c \in D^2$ and hospital h such that c ranks (h, h) and such that h ranks both c_1 and c_2, and for every possible capacity $q \in \{k_h - 2, \ldots, k_h + 2\}$ for h, let $(h_1, h'_1) \ldots, (h_l, h'_l)$ be all the assignments that c prefers to (h, h) and let $D_1, \ldots, D_n$ be all $(q-1)$-tuples of doctors that include neither c_1 nor c_2 and that h prefers to one of c_1 or c_2. We add the following (finite) formula:

$$(\texttt{capacity}_{(h,q)} \wedge \neg(\texttt{matched}_{(c_1,h)} \wedge \texttt{matched}_{(c_2,h)})) \to$$
$$\to \left(\bigvee_{i=1}^{l} (\texttt{matched}_{(c_1,h_i)} \wedge \texttt{matched}_{(c_2,h'_i)}) \vee \bigvee_{i=1}^{n} \bigwedge_{d \in D_i} \texttt{matched}_{(d,h)} \right),$$

which requires that c is not a blocking with (matching both doctors in c to) h.

By construction, and by definition, the models that satisfy all of the above are in one-to-one correspondence with stable matchings between D and H with capacity vectors k^* that differ coordinate-wise from k by at most 2. (The only subtle part is noting that if for some couple c and for some hospital h the couple c does not rank $(h, \emptyset)$ then since the preferences of c are downward closed, this means that no (h, h') are ranked by c and so we have added the formula $\neg\texttt{matched}_{(c_1,h)}$ so indeed it is impossible that c be matched with $(h, \emptyset)$, and symmetrically for when c does not rank $(\emptyset, h)$.) So, it is enough to show that our entire set of formulae is satisfiable, and by the Compactness Theorem, it is enough to show that every finite subset thereof is satisfiable. As a finite set thereof contains only finitely many doctors and finitely many hospitals, this follows from Theorem 4.5, analogously to our proof of Theorem 3.2. □

C Nash Equilibria in Games on Infinite Graphs

In this appendix, we turn to a different setting—games on graphs (see, e.g., Kearns, 2007, and the references therein)—and use Logical Compactness to show the existence of Nash equilibria. We obtain an existence result for games on infinite graphs. Our result here is implicitly covered by Peleg (1969) (who directly generalized the seminal existence result of Nash, 1951), but we give a new proof that uses the same principled approach we use throughout this paper. Like in Section 4.3, we first use Logical Compactness to show the existence of arbitrarily good approximate Walrasian equilibria, and then show that the existence of such approximate Walrasian equilibria implies the existence of exact Walrasian equilibria.

In a *game on a graph*, there is a (potentially infinite) set of players I, each having a finite set of pure *strategies* S_i. Each player $i \in I$ is linked to a finite set of neighbors $N(i) \subset I$ with

$i \in N(i)$, and her utility only depends on the strategies played by players in the set $N(i)$.[34] This setting occurs, for example, in infinite-horizon overlapping-generations models, where at each point in time there are only finitely many players alive, and a player's utility depends only on the behavior of contemporary players. For any player i we denote by $\Sigma_i := \Delta(S_i)$ the set of *mixed strategies* (i.e., distributions over pure strategies) of player i. A *mixed-strategy profile* $(\sigma_i)_{i\in I}$ is a specification of a mixed strategy $\sigma_i \in \Sigma_i$ for every player $i \in I$. A mixed-strategy profile $(\sigma_i)_{i\in I}$ is a *Nash equilibrium* if for every $i \in I$ and every possible deviating strategy $\sigma'_i \in \Sigma_i$, it holds that $u_i(\sigma_{N(i)}) \geq u_i(\sigma'_i, \sigma_{N(i)\setminus\{i\}})$.

Games on *finite* graphs have finitely many players and finitely many strategies per player; hence, the seminal analysis of Nash (1951) implies that they have Nash equilibria.

Theorem C.1 (Follows from Nash, 1951)**.** *Every game on a finite graph has a Nash equilibrium.*

Our main result of this appendix is that Nash equilibria are guaranteed to exist even in games on infinite graphs.

Theorem C.2. *Every game on a (possibly infinite) graph has a Nash equilibrium.*

As already noted, we prove Theorem C.2 by first using the Compactness Theorem to prove the existence of arbitrarily good approximate Nash equilibria, and then showing that the existence of such approximate Nash equilibria implies Theorem C.2. For a given $\varepsilon > 0$, a mixed-strategy profile $(\sigma_i)_{i\in I}$ is an *ε-Nash equilibrium* if for every $i \in I$ and every possible deviating strategy $\sigma'_i \in \Sigma_i$, it holds that $u_i(\sigma_{N(i)}) \geq u_i(\sigma'_i, \sigma_{N(i)\setminus\{i\}}) - \varepsilon$.

Lemma C.3. *For any $\varepsilon > 0$, every (possibly infinite) game on a graph has an ε-Nash equilibrium.*

Proof. Let $\varepsilon > 0$. For each player $i \in I$, the space of profiles of mixed-strategies of players in $N(i)$ is a compact metric space. Specifically, for this proof it will be convenient to consider the space of profiles of mixed-strategies as a metric space with respect to the ℓ^∞ metric;[35] as each player i has a continuous utility function whose domain is this compact metric space, players' utility functions are uniformly continuous by the Heine–Cantor theorem. Thus, there exists $\hat{\delta}_i > 0$ that assures that if two profiles of mixed strategies of players in $N(i)$ are less than $\hat{\delta}_i$ apart, then the utilities they yield to i differs by no more than $\varepsilon/2$.

For each player i, choose $\delta_i := \min\{\hat{\delta}_j \mid j \in N(i)\} > 0$. Recall that Σ_i denotes the space of player i's mixed strategies, and let $\Sigma_i^{\delta_i} \subset \Sigma_i$ be a finite set of strategies that includes all of i's pure strategies, and includes for any mixed strategy in Σ_i a strategy that is at most δ_i away from it; such a set exists by the compactness of Σ_i. For every player i and every profile $\sigma_{N(i)\setminus\{i\}}$ of mixed-strategies for $N(i) \setminus \{i\}$, we define the set of ε-best responses of i:

$$\mathrm{BR}_i^\varepsilon(\sigma_{N(i)\setminus\{i\}}) := \Big\{\sigma_i \;\Big|\; u_i(\sigma_i, \sigma_{N(i)\setminus\{i\}}) \geq \max_{\sigma'_i\in\Sigma_i}\big\{u_i(\sigma'_i, \sigma_{N(i)\setminus\{i\}})\big\} - \varepsilon\Big\}.$$

We now define a set V of variables, and a set of formulae Φ over those variables such that the models that satisfy Φ are in one-to-one correspondence with ε-Nash equilibria where each

[34] Readers familiar with Peleg (1969) will note that even on graphs, Peleg's assumptions are weaker than those stated here. Our analysis can be extended to cover such weaker assumptions, and that our assumptions in other sections can also be similarly weakened. Nonetheless, in general, throughout in this paper we prefer ease and clarity of exposition over tightening assumptions (as noted in the Introduction, we consider the results that we present to be minimal working examples), as our goal is to introduce a unified, transparent technique.

[35] By equivalence of all norms on $\mathbb{R}^n$, the space of profiles of mixed-strategies is also compact with respect to the ℓ^∞ metric.

player i's strategy is in $\Sigma_i^{\delta_i}$. For every player i and discretized strategy $\sigma_i \in \Sigma_i^{\delta_i}$ we introduce a variable $\texttt{plays}_{(i,\sigma_i)}$ that will be TRUE in a model if and only if player i plays the strategy σ_i in the approximate Nash equilibrium that corresponds to the model. We define the set of formulae Φ as follows:

1. For every player $i \in I$ we add the formula

$$\bigvee_{\sigma \in \Sigma_i^{\delta_i}} \texttt{plays}_{(i,\sigma)},$$

   requiring that this player plays some (discretized) strategy; this formula is finite because $\Sigma_i^{\delta_i}$ is.

2. For every player $i \in I$ and distinct strategies $\sigma_i, \sigma'_i \in \Sigma_i^{\delta_i}$, we add the following formula:

$$\texttt{plays}_{(i,\sigma_i)} \to \neg\texttt{plays}_{(i,\sigma'_i)},$$

   requiring that the strategy that player i plays be unique.

3. For every player $i \in I$ and for every profile $\sigma = (\sigma_j)_{j \in N(i)\setminus\{i\}} \in \bigtimes_{j \in N(i)\setminus\{i\}} \Sigma_i^{\delta_i}$ of discretized mixed strategies of $N(i) \setminus \{i\}$, we add the following (finite!) formula:

$$\Bigg(\bigwedge_{j \in N(i)\setminus\{i\}} \texttt{plays}_{(j,\sigma_j)}\Bigg) \to \Bigg(\bigvee_{\sigma_i \in \Sigma_i^{\delta_i} \cap \mathrm{BR}_i^{\varepsilon}(\sigma)} \texttt{plays}_{(i,\sigma_i)}\Bigg),$$

   requiring that player i ε-best-responds to the strategies played by the other players.

We claim that the preceding set of formulae is satisfied by some model. To see this, we first note that any finite subset of the formulae mentions only finitely many players. Now, consider the game between those players, where all other "players" mechanically play their first strategy. This restricted game has a Nash equilibrium by Theorem C.1. By choosing for each player a closest strategy in $\Sigma_i^{\delta_i}$, each player's utility changes by at most $\varepsilon/2$ (by uniform continuity), and so does the utility attainable by best responding. Therefore, since we started with a Nash equilibrium, it is assured that each player is now playing an ε-best response, and so the finite subset of formulae is satisfied. Hence, by the Compactness Theorem, the collection of all formulae is satisfied by some model, and thus the game has an ε-Nash equilibrium. $\square$

Now, we can use Lemma C.3 to prove Theorem C.2 by way of a "diagonalization" argument.

Proof of Theorem C.2. Since each player in the graph has finitely many neighbors, every connected component of the graph consists of at most countably many players. As it is enough to show the existence of a Nash equilibrium in each connected component separately (we use the Axiom of Choice here), let us focus on one connected component. By Lemma C.3 there exists a sequence $(\sigma^n)_{n=1}^{\infty}$ of $\frac{1}{n}$-Nash equilibria in the game on this connected component. Since each of the at-most-countably-many coordinates of each element in this sequence lies in $[0,1]$, we can choose a subsequence (a "diagonal subsequence") that converges in all coordinates; let σ^* denote the limit of that subsequence.

We claim that σ^* is a Nash equilibrium. To see this, note that for every $i \in I$ and $\sigma'_i \in \Sigma_i$, we have for the nth elements of the sequence that

$$u_i(\sigma^n_{N(i)}) \geq u_i(\sigma'_i, \sigma^n_{N(i)\setminus\{i\}}) - \tfrac{1}{n}.$$

By the continuity of u_i, (C) means that for every $i \in I$ and $\sigma_i' \in \Sigma_i$, we have

$$u_i(\sigma^*_{N(i)}) \geq u_i(\sigma_i', \sigma^*_{N(i)\setminus\{i\}}),$$

so no player has a profitable deviation under the profile σ^*. Hence, σ^* is indeed a Nash equilibrium—and in particular, we see that a Nash equilibrium exists in the game, as desired. $\square$