HARVARD | BUSINESS | SCHOOL

A Compact, Logical Approach to Large-Market Analysis

Yannai A. Gonczarowski
Scott Duke Kominers
Ran I. Shorrer

Working Paper 19-127

A Compact, Logical Approach to Large-Market Analysis

Yannai A. Gonczarowski
Tel Aviv University

Scott Duke Kominers
Harvard Business School

Ran I. Shorrer
Penn State University

Working Paper 19-127

A Compact, Logical Approach to Large-Market Analysis*

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

June 25, 2019

Abstract

In game theory, we often use infinite models to represent "limit" settings, such as markets with a large number of agents or games with a long time horizon. Yet many game-theoretic models incorporate finiteness assumptions that, while introduced for simplicity, play a real role in the analysis. Here, we show how to extend key results from (finite) models of matching, games on graphs, and trading networks to infinite models by way of *Logical Compactness*, a core result from Propositional Logic. Using Compactness, we prove the existence of man-optimal stable matchings in infinite economies, as well as strategy-proofness of the man-optimal stable matching mechanism. We then use Compactness to eliminate the need for a finite start time in a dynamic matching model. Finally, we use Compactness to prove the existence of both Nash equilibria in infinite games on graphs and Walrasian equilibria in infinite trading networks.

1 Introduction

In game theory, we often think of infinite models as better than finite ones for representing limit behavior: infinite sets of agents, for example, correspond to "large markets," while infinite-horizon dynamic games represent interactions that will be repeated over and over with no fixed start or end time. Yet many game-theoretic models incorporate finiteness assumptions for simplicity—and those assumptions often play a real role in the analysis.

Nevertheless, the core results in game theory are so elegant that it seems they should logically extend to infinite models, as well. In this paper, we show that the preceding intuition is precisely—and in fact, verbatim—correct, at least for a number of matching, graph, and trading network settings. We show, for example, that classic existence, structure, and strategy-proofness results for stable matching in *finite* markets imply analogous results in a large class of infinite-size markets, by

*The authors thank Ravi Jagadeesan and David Laibson for helpful comments. Gonczarowski was supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities; his work was supported by ISF grant 1435/14 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). Additionally, Kominers gratefully acknowledges the support of 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).

†The Blavatnik School of Computer Science, Tel Aviv University, Israel — *E-mail*: yannai@gonch.name. Part of the work of Gonczarowski was carried out while he was at the Hebrew University of Jerusalem and Microsoft Research.

‡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.

way of a result central to the theory of Propositional Logic. An analogous approach lets us prove existence of stable matchings in infinite-horizon dynamic matching markets (with neither a finite "start" nor "end" time), as well as the existence of Nash equilibria in games on infinite graphs and the existence of Walrasian equilibria in infinite trading networks.

The key to our approach is *Logical Compactness*, which (roughly) states that an infinite set of individually finite logical statements can be made consistent if and only if every finite subset can.[1] We show how to rephrase stable matching problems, network games, and economies with indivisible goods in terms of a collection of individually finite logical statements; Logical Compactness then enables us to lift results from finite models to infinite ones.

To encode a (one-to-one) stable matching, for example, we consider a set of Boolean variables $\{\texttt{matched}_{(m,w)}\}$ that are TRUE when man m is matched to woman w. 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 cardinality 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 we introduce some mild assumptions on agents' preferences, 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).[2]

The same idea lets us generalize man-optimality and strategy-proofness results for stable matching to infinite markets, as well—although the arguments required to apply Logical Compactness in those cases are more subtle.[3] Likewise, we can use Logical Compactness to extend matching models to incorporate infinite horizons: We consider a dynamic matching setting inspired by 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 iteratively; 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 (see, e.g., Öry, 2016; Clark, Fudenberg, and Wolitzky, 2019) as a representation of a steady state.[4]

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

[2] A similar approach can be used to generalize the matching with contracts setting of Kelso and Crawford (1982), Hatfield and Milgrom (2005), and Hatfield and Kominers (2017) to infinite settings.

[3] As we discuss in the sequel, existence and man-optimality of stable matchings in the infinite settings we consider were originally proven by Fleiner (2003). Our strategy-proofness result is original to the present work. And indeed, standard proofs of strategy-proofness rely on versions of the Lone Wolf/Rural Hospitals Theorem (Roth, 1984), which Jagadeesan (2018b) has shown fails in a setting that is a special case of our model.

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

And Logical Compactness has microeconomic applications beyond matching: among other results, we use Logical Compactness to obtain existence results for both Nash equilibria in games on graphs as well as Walrasian equilibria in trading networks. In the case of both Nash and Walrasian equilibria (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.

Overall, our methods give a new way of showing the existence of—and characterizing—limit objects in game theory, and in economic theory, more broadly. Moreover, in the limits we obtain under Logical Compactness, agents in a sense "maintain their mass," which may be attractive for economic applications because strategic issues survive the limiting process.

Working with Logical Compactness reduces reasoning about infinite problems to reasoning about finite problems—and we can think about those finite problems directly without (say) having to consider limiting processes and/or convergence. Indeed, once we have a logical formulation, compactness arguments permit us to reason solely about finite models in their regular language of analysis. 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.

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. We therefore view our contribution first and foremost as introducing a unifying approach that is simple and intuitive to work with, and does not require us to look for the "right" topological space or apply topological reasoning directly.[5]

In each setting we consider, we give what can be thought of as minimal working examples; it seems likely that Compactness can be used to further generalize results in these and other settings, as well (even before translating into topological statements as discussed above). For example, we suspect that the approach we use in Section 4 should translate directly to other dynamic matching models (e.g., Doval, 2017; Kadam and Kotowski, 2018), including those with infinitely many periods. It seems logical that the results we present represent just the tip of the iceberg in terms of what the techniques we introduce—and future refinements thereof—should eventually be able to accomplish.

1.1 Related Literature

Infinite models are used frequently in game/economic theory as a way of representing limit—or "large"—markets. Some "large market" models feature 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, 2019; Greinecker and Kah, 2019). A second class of models has worked 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).[6] Our work is much more closely related to models of the second type, as we require some amount of local countability in most of the settings we consider.

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

[5] Once a proof is derived using our approach, it may be possible to translate it to a topological statement and attempt to achieve greater generality, if such generality is of interest.

[6] A series of recent papers has introduced models that mix between these 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).

characterization of stable matchings that is independent of the cardinality of the set of agents. 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) that requires that each agent find at most finitely many partners acceptable. Our results for games on graphs (see, e.g., Kearns, 2007, and the references therein) are implicitly covered by Peleg (1969), who directly generalizes the seminal existence result of Nash (1951). But while some of our matching and Nash equilibrium results are implicitly covered by previous analyses (Fleiner, 2003; Peleg, 1969), their respective known proofs rely on a different, highly specialized arguments. By contrast, our proofs of these results—and furthermore our proofs of all of the other 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. But at the same time, we use our methods to prove completely new results: our results for trading networks—all of which are novel to this work—generalize the existence result of Hatfield et al. (2013) to infinite trading networks (see also Hatfield et al., 2019, 2018; Fleiner et al., forthcoming). Meanwhile, we show a way that Logical Compactness can be used to convert a dynamic game with a finite start time into an "ongoing" dynamic game with neither start nor end—connecting 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 Logical Compactness to reinterpret and generalize results in economics. That said, 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).[7]

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 continuous objects (such as Nash equilibria and Walrasian prices), as well as complex characterization results (e.g., strategy-proofness) that go far beyond existence results and are not inherently topological.

In addition to the fact that Hellman and Levy (2019) use (different) tools from mathematical logic, their paper is also conceptually related to ours. In a precise sense, their paper and ours complement each other: our paper gives a principled approach to lifting existence results (and beyond) from finite markets to countable ones, while Hellman and Levy (2019) give a principled approach to lifting existence results from infinitely-countable (but not from finite) markets to uncountable ones. Combining both approaches enables us to lift various existence results from finite to uncountable markets.

1.2 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 first

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

gives another simple illustration of our approach—using Logical Compactness to (re-)prove the existence of stable outcomes in infinite, one-to-one matching markets—and then states our structural and strategy-proofness results, deferring their proofs to Section 7. Section 4 uses Logical Compactness to generalize the dynamic matching framework inspired by work of Pereyra (2013). Section 5 uses Logical Compactness to prove the existence of Nash equilibria in games on graphs. Section 6 then uses Logical Compactness to prove the existence of Walrasian equilibria in infinite trading networks. Section 7 uses Logical Compactness to prove the structural and strategy-proofness results introduced in Section 3. Section 8 concludes.

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.[8] 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 logic model; 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:*

$$\text{`}\mathtt{P}\text{'}, \tag{1}$$

$$\text{`}(\mathtt{P} \vee \mathtt{Q})\text{'}, \tag{2}$$

$$\text{`}\neg(\mathtt{P} \wedge \mathtt{Q})\text{'}, \tag{3}$$

$$\text{`}((\mathtt{P} \wedge \mathtt{R}) \to \mathtt{S})\text{'}. \tag{4}$$

[8] 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).

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 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 $X \subseteq \mathbb{N}$,*

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

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;
- '$(\phi \rightarrow \psi)$' is TRUE iff either ϕ is FALSE or ψ is TRUE or both (that is, '$(\phi \rightarrow \psi)$' 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:*

'$\mathtt{P}$',	*"$\mathtt{P}$ [IS TRUE]"*	(1)
'$(\mathtt{P} \vee \mathtt{Q})$'	*"$\mathtt{P}$ OR $\mathtt{Q}$ [IS TRUE]",*	(2)
'$\neg(\mathtt{P} \wedge \mathtt{Q})$'	*"NOT ($\mathtt{P}$ AND $\mathtt{Q}$ [ARE BOTH TRUE])",*	(3)
'$((\mathtt{P} \wedge \mathtt{R}) \rightarrow \mathtt{S})$'	*"$\mathtt{P}$ AND $\mathtt{R}$ [BOTH BEING TRUE], IMPLIES $\mathtt{S}$".*	(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). We suspect that the proof we present here may be more accessible to quite a few students. Indeed, Logical Compactness, like Zorn's Lemma, relies on the Axiom of Choice—but parsing the statement of the Compactness Theorem, as well as applying it, may require less background.[9]

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.

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 $>_X|_{X'}$; the model corresponding to that order satisfies Φ' (e.g., Lahiri, 2002).[10] 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. □

Our approach in this paper is to 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

[9] Mas-Colell, Whinston, and Green (1995) label their proof (which uses Zorn's Lemma) as "advanced."

[10] 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.

(in the same sense) to a finite variant of the infinite problem. Thus, known results for solutions 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 Stable Matching in Infinite Markets

3.1 Setting

We work with 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[11]) 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.

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.[12]

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

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 $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$.

[11]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.

[12]For alternative proofs via a fixed-point argument, or via an infinite variant of Gale and Shapley's algorithm, see Fleiner (2003) and Jagadeesan (2018b), respectively.

- For every man m and for every two women $w \neq w'$, we add the following formula to $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$:
$$\texttt{matched}_{(m,w)} \to \neg\texttt{matched}_{(m,w')}, \tag{5}$$
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 $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$:
$$\texttt{matched}_{(m,w)} \to \neg\texttt{matched}_{(m',w)}, \tag{6}$$
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.[13]) We add the following (finite!) formula to $\Phi_{(M,W,\mathcal{P}_M,\mathcal{P}_W)}$:
$$\neg\texttt{matched}_{(m,w)} \to \Big(\big(\texttt{matched}_{(m,w_1)} \vee \cdots \vee \texttt{matched}_{(m,w_l)}\big) \vee \vee \big(\texttt{matched}_{(m_1,w)} \vee \cdots \vee \texttt{matched}_{(m_k,w)}\big)\Big), \tag{8}$$
which effectively requires 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 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 $(M, W, \mathcal{P}_M, \mathcal{P}_W)$. □

[13]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 richer types of infinite preference lists, under which such a construction would not be possible.

As described in the Introduction, once we translate the stable matching problem into a series of locally finite logical statements (here, statements about variables specifying who matches with whom), Logical Compactness lets us take the classical existence result for finite markets (Theorem 3.1) and lift it to infinite markets (Theorem 3.2).

While one can prove Theorem 3.2 using other methods (again, see Fleiner, 2003; Jagadeesan, 2018b), as we explain next, Compactness lets us extend structural and strategic results for matching, as well.

3.3 Structure and Incentives

In contrast to our nonconstructive proof of Theorem 3.2, Gale and Shapley's the 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 3.3 (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 3.3 to infinite markets. Indeed, while Theorem 3.3 is also an existence result, it is a far more intricate one, which may be thought of as a "second-order" existence result: 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 3.3 to infinite markets by way of Logical Compactness, as well—but the argument is far more intricate than our proof of Theorem 3.2.[14]

Theorem 3.4. *In any (possibly infinite) one-to-one matching market, there exists a man-optimal stable matching.*

With Theorem 3.4 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 these 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 3.5 (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 3.5 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

[14] Again, Theorem 3.4 was originally proven by Fleiner, 2003.

Wolf/Rural Hospitals Theorem (Roth, 1984), which Jagadeesan (2018b) has shown does not hold in our setting (or even in less general environments), so a significant innovation on the proof strategy is needed here—and moreover, the mere applicability of Theorem 3.5 to infinite markets is not *a priori* clear. Third, Theorem 3.5 is of a very different flavor than results that have traditionally been lifted to infinite settings using Logical Compactness. Indeed, it is not an existence result, 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 called Theorem 3.4 a "second-order" existence result, then Theorem 3.5 is an "at least third-order" result that is quite far from an 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 a claim to which the strategy-proofness of the man-optimal stable mechanism can be carefully reduced, and hence to prove that this mechanism is strategy-proof.

Theorem 3.6. *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.*

We emphasize that unlike Theorems 3.2 and 3.4, Theorem 3.6 is novel to the present work—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 3.4, Theorem 3.6 is not an existence result about matchings with certain properties—it deals with the incentive properties of a specific matching mechanism. Thus, Theorem 3.6 illustrates that Logical Compactness has applications beyond results we might naturally expect to be able to prove with limiting or continuity arguments.

Our proofs of Theorems 3.4 and 3.6 are more involved: 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. Thus, we defer the proofs of Theorems 3.4 and 3.6 to Section 7, by which point we will have built up additional methodological intuition about how to use Logical Compactness in game-theoretic settings.

4 Stable Matching with an 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."

4.1 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 $t-1$.[15]

4.2 Challenge

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 above argument 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 4.1. 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.

[15] 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).

- 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 4.1? 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.

4.3 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).

Definition 4.2. 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 4.3. *In any (possibly infinite) dynamic matching market that has finite presence, a stable-subject-to-tenure matching chronology exists.*

Before proving Theorem 4.3, we note that the finite presence condition in Theorem 4.3 cannot be dropped. Indeed, 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 4.3. 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.

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

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

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

- 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.

- 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.

- 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 the existence result of Pereyra (2013) using the minimum involved time as a "period 0." □

5 Nash Equilibria in Games on Infinite Graphs

In this section, 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. The application of Compactness here is more complex than that used in the preceding sections: we first use the Compactness Theorem to show the existence of arbitrarily close approximations of Nash equilibria in the infinite graph, and then show that the existence of approximate Nash equilibria implies the existence of exact Nash equilibria.

5.1 Setting

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 $N(i)$.[16] 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 $u_i(\sigma_{N(i)}) \geq u_i(\sigma'_i, \sigma_{N(i)\setminus\{i\}})$.

5.2 Existence

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

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

Theorem 5.2. *Every game on a (possibly infinite) graph has a Nash equilibrium.*

As already noted, we prove Theorem 5.2 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 5.2. 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 5.3. *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;[17] 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:

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

[16] 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.

[17] 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.

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 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:

- 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.

- 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.

- 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 5.1. 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. □

Now, we can use Lemma 5.3 to prove Theorem 5.2 by way of a "diagonalization" argument.

Proof of Theorem 5.2. 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[18]), let us focus on one connected component. By Lemma 5.3 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

[18] A better-behaved way to choose a Nash equilibrium for each connected component (once we have proven that such equilibria exist) is to use the result of Hellman and Levy (2019), which guarantees that the overall equilibrium be measurable. (See also the discussion of that paper in Section 1.1.)

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, (5.2) 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. □

6 Walrasian Equilibria in Infinite Trading Networks

We next turn to an infinite variant of the trading-network matching framework of Hatfield et al. (2013) and show the existence of Walrasian equilibria. Like in Section 5, 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.

6.1 Setting

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, $\Omega_i := \big\{\omega \in \Omega \;\big|\; i \in \{\mathsf{s}(\omega), \mathsf{b}(\omega)\}\big\}$. 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 quasi-linear in prices and otherwise depends only on the set of trades $\Omega'_i \subseteq \Omega_i$ that are executed.

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 $O_i := \big\{o \in O \;\big|\; i \in \{\mathsf{s}(\mathsf{t}(o)), \mathsf{b}(\mathsf{t}(o))\}\big\}$ 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 $\big\{\mathsf{o}(\omega) \mid \omega \in \Omega'_i \text{ and } \mathsf{b}(\omega) = i\big\} \cup \big\{\mathsf{o}(\omega) \mid \omega \notin \Omega'_i \text{ and } \mathsf{s}(\omega) = i\big\}$; 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, $u_i\big(\{\mathsf{o}(\omega) \mid \mathsf{s}(\omega) = i\}\big) = 0$; 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.

Definition 6.1 (Gross Substitutability)**.** The preferences of agent i are *(gross) substitutable* if for all price vectors $p, p' \in \mathbb{R}^O$ such that $\big|D_i(p)\big| = \big|D_i(p')\big| = 1$ and $p \leq p'$, if $o \in D_i(p)$ then $o \in D_i(p')$ for each $o \in O_i$ such that $p_o = p'_o$.

Definition 6.2 (Walrasian Equilibrium)**.** 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 $i \in I$.

6.2 Existence

Hatfield et al. (2013) have shown that substitutable preferences suffice to guarantee the existence of Walrasian equilibria in finite trading networks.

Theorem 6.3 (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 6.3 to infinite trading networks.

Theorem 6.4. *If all agents have substitutable preferences, then every (possibly infinite) trading network has a Walrasian equilibrium.*

As already noted, we prove Theorem 6.4 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 6.4. We start by defining precisely what we mean by approximate Walrasian equilibria.

Definition 6.5 (Approximate Walrasian Equilibrium)**.** 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 6.6. *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 $\big|u_i(O'_i \cup \{o\})\big| + \big|u_i(O'_i)\big| < \infty$ then $\big|u_i(O'_i \cup \{o\}) - u_i(O'_i)\big| < H_o$. (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:

- 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.

- 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.

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

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

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

- 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 6.3 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 6.6 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 6.6 nevertheless yields the existence of an exact Walrasian equilibrium.

Proof of Theorem 6.4. 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[19]), let us focus on one connected component. For a diminishing sequence $\varepsilon_n \to 0^+$, by Lemma 6.6 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. □

7 Proofs of Theorems 3.4 and 3.6

Having built up more intuition about how use Logical Compactness for game-theoretic/economic problems, we now return to the setting of Section 3 and prove Theorems 3.4 and 3.6.

We start by proving Theorem 3.4. As already noted above, the proofs that we present below are more involved than those presented so far, and this already manifests in the proof of Theorem 3.4, 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; and
- the argument invokes additional results from the literature on finite matching markets.

Proof of Theorem 3.4. 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.

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

[19] Once again, a better-behaved way to choose a Walrasian equilibrium for each connected component (once we have proven that such equilibria exist) is to use the result of Hellman and Levy (2019), which guarantees that the overall equilibrium be measurable. (See also the discussion of that paper in Section 1.1.)

$m \in \widehat{M}$, 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).[20]

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 3.4 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 3.3 it has a man-optimal stable matching—we will show that this matching satisfies all of the formulae in the finite subset of formulae.[21]

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.

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. □

[20] 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.

[21] 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').

Finally, we turn to proving Theorem 3.6. The argument we present has all of the complexities described previously, along with one more:

- We have to find a way to reduce Theorem 3.6—which is quite far from an existence result—to an existence result to which we can apply the Compactness Theorem.

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

Proof of Theorem 3.6. By way of contradiction, assume that there exist profiles of preferences $\mathcal{P}_M = (P_m)_{m \in M}$ and $\mathcal{P}_W$, and some man $\tilde{m}$ who can profitably manipulate the man-optimal mechanism given this preference profile by misreporting his preferences to be $P'_{\tilde{m}}$ rather than $P_{\tilde{m}}$. Then, by the individual rationality property of stable matchings and the fact that the manipulation is successful, the manipulation leads $\tilde{m}$ to be matched to some woman $\tilde{w}$. From the definition of the man-optimal mechanism, it must be that $\tilde{m}$ strictly prefers $\tilde{w}$ to any woman to whom he is matched in any matching that is stable with respect to $\mathcal{P}_M$ and $\mathcal{P}_W$.

We first claim that by reporting to the mechanism the preference list 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 $(M, W, (\{\tilde{w}\}, \mathcal{P}_{M\setminus\{\tilde{m}\}}), \mathcal{P}_W)$.

Next, we claim that $\tilde{m}$ will also be matched with $\tilde{w}$ or a better-preferred woman by reporting his true preferences truncated at $\tilde{w}$ (i.e., reporting only the women he likes at least as much as $\tilde{w}$ in his true order of preference). For this, we use Logical Compactness.[22]

We denote by $P^{\tilde{w}}_{\tilde{m}}$ the truncation of $\tilde{m}$'s true preferences $P_{\tilde{m}}$ at $\tilde{w}$. We work with the variables $V := V_{(M,W,(P^{\tilde{w}}_{\tilde{m}},\mathcal{P}_{M\setminus\{\tilde{m}\}}),\mathcal{P}_W)}$ as defined in the proof of Theorem 3.2, and as the set of formulae Φ we take the formulae $\Phi_{(M,W,(P^{\tilde{w}}_{\tilde{m}},\mathcal{P}_{M\setminus\{\tilde{m}\}}),\mathcal{P}_W)}$ from that proof and add to them the following (finite!) formula:

$$\bigvee_{w \succcurlyeq_{P^{\tilde{w}}_{\tilde{m}}} \tilde{w}} \texttt{matched}_{(\tilde{m},w)}$$

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 list is $P^{\tilde{w}}_{\tilde{m}}$ 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 3.5, it must be that in the same finite market, the man-optimal stable matching with respect to $P^{\tilde{w}}_{\tilde{m}}$ has $\tilde{m}$ matched to a woman he $P^{\tilde{w}}_{\tilde{m}}$-ranks at

[22] As promised, we 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.

least as high as $\tilde{w}$ (otherwise, had $\tilde{m}$'s true preferences been $P^{\tilde{w}}_{\tilde{m}}$, 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 $P^{\tilde{w}}_{\tilde{m}}$). Hence the entire collection of formulae is satisfied by some model, by the Compactness Theorem.

But a model that satisfies the entire set of formulae corresponds to a stable matching with respect to the true profile of preferences $(\mathcal{P}_M, \mathcal{P}_W)$. This is because $\tilde{m}$ is matched to $\tilde{w}$ or a woman he prefers more, and we only removed from his list women he likes less than $\tilde{w}$, and so we did not remove any opportunities to block the match. Hence, there exists a stable matching with respect to the true preferences that matches $\tilde{m}$ to a woman he likes more than the most preferred woman to whom he is matched in any stable matching—a contradiction. □

8 Conclusion

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

As we have demonstrated, the Compactness Theorem for Propositional Logic can be used to analyze large-market models of matching, games on graphs, and exchange economies, and can also be used to extend repeated games with finite start-times into "ongoing" dynamic games with neither start nor end. While the Compactness Theorem is implied by more general results in topology, we can use it to prove results (such as strategy-proofness) that are not topological in nature. Moreover, working with propositional logic lets us reason about infinite models without having to formulate explicit topological arguments.

A distinguishing feature of the large-market equilibria we obtain using Compactness is that agents in those limits "maintain their mass," in the sense that they are still subject to the strategic incentives present in the associated finite models. It is thus possible that Compactness-based approaches may yield especially realistic large-market models of economic behavior. Our limits may also be viewed as part of a path towards limits with uncountably many agents—as we mentioned in the Introduction, combining our results with those of Hellman and Levy (2019) makes it possible to lift existence results from finite markets to uncountable ones.

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, without having to consider limiting processes and/or convergence. We hope this will make Compactness both easily and flexibly applicable in a range of economic theory contexts. We might even one day hope to use Compactness to develop a method of translating results for a large class of finite (or finite-horizon) markets to infinite ones that could be applied without need for setting-specific infrastructure.

References

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

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 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*, 2019. forthcoming.

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

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

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

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

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

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. Working Paper, 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.

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

T. Fleiner, R. Jagadeesan, Z. Jankó, and A. Teytelboym. Trading networks with frictions. *Econometrica*, forthcoming.

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*, 2019. forthcoming.

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. Working paper, 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.

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. W. Hatfield and S. D. Kominers. Contract design and stability in many-to-many matching. *Games and Economic Behavior*, 101:78–97, 2017.

J. W. Hatfield and P. Milgrom. Matching with contracts. *American Economic Review*, 95(4): 913–935, 2005.

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. Chain stability in trading networks. Stanford University Working Paper, 2018.

J. W. Hatfield, S. D. Kominers, A. Nichifor, M. Ostrovsky, and A. Westkamp. Full substitutability. *Theoretical Economics*, 2019. forthcoming.

Z. Hellman and Y. J. Levy. Measurable selection for purely atomic games. *Econometrica*, 2019. forthcoming.

N. Immorlica and M. Mahdian. Marriage, honesty, and stability. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 53–62. Society for Industrial and Applied Mathematics, 2005.

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

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.

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.

A. S. Kelso, Jr. and V. P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.

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

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

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. Working Paper, 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.

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

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

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

M. K. Richter. Revealed preference theory. *Econometrica (pre-1986)*, 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.

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