EVERY OBSTRUCTION TO NONCONTEXTUAL HIDDEN VARIABLES IS FINITE: WELL-FOUNDED REFUTATIONS IN INFINITARY REAL-VALUED LOGIC

OLIVIA AUBONE AND JOSÉ IOVINO

Abstract. We study the obstructions to noncontextual hidden variables when quantum probability is described in full, by the complete lattice of closed subspaces of a separable infinite-dimensional Hilbert space together with its countably additive measures. We give such a description in infinitary real-valued logic, the infinitary extension of Keisler's model theory of general structures; the countable suprema and infima of this logic express the completeness of the lattice and the countable additivity of the measures. The standard models of our axioms represent exactly the density-operator measures of Gleason's theorem. The axioms are closed conditions, non-strict inequalities between the values of sentences and rational numbers. The expressive power afforded by infinitary sentences comes at the price of compactness. The hypothesis that at each hidden state every event receives the value 0 or 1, whatever the context of measurement, contradicts the axioms, yet it is consistent with every finite subset of them. The obstructions nevertheless remain finite. We prove that in this logic a countable set of closed conditions is inconsistent if and only if a well-founded tree whose branches end in finite contradictions refutes it; for the hypothesis above, every such tree has an infinitely branching node.

The finiteness theorem is therefore maximal: a hidden-variable hypothesis expressed by countably many closed conditions is inconsistent with the axioms, and thus an obstruction, exactly when such a tree refutes the hypothesis together with the axioms. The same language defines, uniformly across standard models, exactly the σ-projective sets of probabilities, and all these results are proved in ZFC.

Overview

We study the obstructions to noncontextual hidden variables when quantum probability is described in full. An outcome-deterministic noncontextual hidden-variable model assigns to every event the value 0 or 1 at each hidden state, whatever the context in which the event is measured. An obstruction is a hypothesis about hidden variables, expressed by countably many closed conditions, that has no model together with the axioms of quantum probability, standard or not. The full description is the setting of Gleason's theorem: the complete lattice of closed subspaces of a separable infinite-dimensional Hilbert space, with its countably additive probability measures.

Our description is infinitary. In infinite dimension the lattice of closed subspaces is complete and probability is countably additive, and our axioms in infinitary real-valued logic describe both through countable suprema and infima and through a sort

2020 *Mathematics Subject Classification.* Primary 03C75; Secondary 81P13, 81P10, 03C66, 03C25, 91A44, 03E15, 03G12, 06C15.

Key words and phrases. contextuality, noncontextual hidden variables, Kochen–Specker theorem, Gleason's theorem, countable additivity, failure of compactness, well-founded refutation trees, infinitary logic, real-valued logic, model existence game, topological games, model-theoretic forcing, uniform definability, σ-projective sets, quantum logic.

This version: November 13, 2026. The most recent version of this paper is available at https://utsa-quantum-math.github.io/resources/gleason-kochen-specker-game/; it may supersede the version deposited on arXiv.

of square-summable sequences of vectors. The logic is interpreted in Keisler's general structures, which prescribe neither a metric nor a modulus of uniform continuity; the lattice operations can therefore be interpreted as they are, although the meet of two distinct rays is the zero subspace however close the rays are. In the standard models, those built on full Hilbert spaces, the complete lattice is recovered and the axioms single out exactly the density-operator measures of Gleason's theorem; the axioms also have nonstandard models, among them countable ones.

The expressive power of infinitary sentences come at the price of compactness. The dispersion-free hypothesis, that every probability is 0 or 1, is a single closed condition. Adjoined to the axioms, it yields a theory with no models, by a proof that uses neither Gleason's theorem nor standardness, and yet every finite subset of the axioms of that theory has a model.

The obstructions nevertheless remain finite. Our refutation theorem applies to countable theories of closed conditions, whose axioms bound the value of a sentence above or below by a rational number: such a theory is inconsistent exactly when it is refuted by a well-founded tree whose leaves are finite sets of interval commitments that no structure satisfies together. The commitments concern sentences that may themselves be infinitary, and for the axioms as presented the finiteness resides in the commitments rather than in finite subsets of the axioms. The trees are winning strategies in a topological game of model existence, which open determinacy decides; a tree branches infinitely only at countable suprema and infima, and for the dispersion-free theory every refuting tree has such a node.

The resulting finiteness theorem is maximal. A hidden-variable hypothesis given by countably many closed conditions has no model together with the axioms, standard or not, if and only if such a tree refutes the two together; the certified hypotheses are therefore exactly the obstructions, and no sound method of refutation certifies more.

The refutations are explicit. In the refutation we construct for the dispersion-free hypothesis, countable additivity compels a commitment that some ray has probability close to 1, and the ray selects a four-dimensional subspace in which a parity argument closes every branch. Geometric axioms identify the occurrences of a ray in different contexts, while a separate measure axiom carries the probability across that identity; the refutation thus isolates each use of noncontextuality, and its contradiction survives explicit numerical tolerances.

The earlier logical formulations of contextuality derive their finite results from compactness, in finite dimension. In J. L. Bell's propositional formulation, which records the values 0 and 1 of a noncontextual assignment, compactness extracts from Gleason's theorem the finite configurations of Kochen and Specker in each finite dimension at least three. Hrushovski and Pitowsky, working in classical two-valued first-order logic with probabilities as elements of a sort of scalars, obtain from compactness and the decidability of real closed and algebraically closed fields effective finite approximations to the predictions of a quantum state. The infinite-dimensional setting of Gleason's theorem lies beyond propositional and first-order logic, and compactness is no longer available in it. The refutation theorem takes its place, and by the maximality described above its well-founded trees of finite contradictions certify every obstruction.

The expressive power of our infinitary language, which describes the complete lattice together with its countably additive measures, also determines which sets of probabilities are definable: uniformly across standard models, the sets definable through countable suprema and infima and quantification over events are exactly the σ-projective subsets of the unit interval. All our results are proved in ZFC.

Contents

Introduction 3
Main results of the paper 7
Organization of the paper 10
1. The geometric setting 12
2. The logic 13
2.1. Real-valued $L_{\omega_1\omega}$ 13
2.2. Henkin types and model existence 15
3. The model existence game and finite obstructions 17
3.1. The setting 17
3.2. The model existence theorem as a game 21
3.3. Finite obstructions 23
4. The vocabulary 27
5. The axioms 30
6. Standard models of T_0: pre-Solèr lattices 36
7. Standard models of T_1 and T: Gleason's theorem, the range of the measure and definable sets of values 39
7.1. Definable sets of values 43
8. The proof of Theorem C 48
8.1. The Bell–Kochen–Specker obstruction in the game 52
9. Finite obstructions record what a Kochen–Specker configuration assumes 58
9.1. Noncontextuality and determinism concern the responses at a single state 59
9.2. Countable additivity chooses the configuration 60
9.3. The game records noncontextuality as the substitutivity of equality 61
9.4. The parity obstruction survives numerical tolerances 64
10. Remarks 65
Appendix A. Lattices, hermitian spaces and projective geometry 68
A.1. Lattices 69
A.2. Hermitian spaces 71
A.3. Projective geometry, Solèr lattices and pre-Solèr lattices 73
A.4. What a pre-Solèr lattice is missing 80
A.5. Nonexpressibility over pre-Solèr lattices 82
Appendix B. The space of Henkin types and the model existence game 83
The space of Henkin types 83
Positions and plays 84
Closed conditions, strict conditions and compactness 86
Baire category 86
References 88

Introduction

We study the obstructions to noncontextual hidden variables when quantum probability is described in full, by the complete lattice of closed subspaces of a separable infinite-dimensional Hilbert space together with its countably additive measures. By an obstruction we mean a hypothesis about hidden variables that has no model together with our axioms for quantum

probability, standard or nonstandard[1]. We show that every obstruction is finite: a well-founded tree refutes the hypothesis together with the axioms, and the terminal contradictions of the tree are finite sets of commitments.

The mathematical setting is that of Gleason's theorem. For a Hilbert space H, the closed subspaces form a lattice $L(H)$, and each closed subspace is an event, identified with the orthogonal projection onto it. A density operator on H is a positive operator of trace 1 and represents a quantum state. When H is separable and of dimension at least three, Gleason's theorem [Gle57] identifies the countably additive probability measures on $L(H)$ with the maps $P \mapsto \operatorname{Tr}(\rho P)$ given by density operators ρ, none of which takes only the values 0 and 1. An outcome-deterministic noncontextual hidden-variable model assigns to every event, at each hidden state, a value 0 or 1; in every complete set of mutually orthogonal events exactly one event receives the value 1, and the value of an event does not depend on which compatible events are measured together with it [BCG$^+$22]. Specker announced the impossibility of such assignments in dimension three [Spe60]. J. S. Bell observed that such an assignment is a $\{0,1\}$-valued measure on $L(H)$, which Gleason's theorem excludes [Bel66]. Kochen and Specker then exhibited a finite configuration of rays that no such assignment respects [KS67].

We work in infinitary real-valued logic [ADI26], which we denote $L_{\omega_1\omega}(\mathbb{R})$: the infinitary extension of Keisler's model theory of general structures [Kei23, Section 2.1]. In this logic the truth values of formulas are real numbers, and the probability of an event is the value of a predicate. The countable suprema and infima of $L_{\omega_1\omega}(\mathbb{R})$ express two features of the full setting: the completeness of the lattice of closed subspaces and the countable additivity of probability. Completeness is not expressible in first-order logic, and we show that among pre-Solèr lattices it is not expressible by any sentence of classical two-valued $L_{\omega_1\omega}$ in the language of ortholattices expanded by constants for countably many atoms (Proposition A.24). Countable additivity is not implied by finite additivity: on the finite-dimensional subspaces of an infinite-dimensional space and their orthogonal complements, the assignment of 0 to every finite-dimensional subspace and of 1 to every complement is a finitely additive probability measure (Proposition 9.1), which only countable additivity excludes. In general structures, a vocabulary assigns to each predicate symbol an a priori bound, and a structure interprets function and predicate symbols by arbitrary maps that respect these bounds. There is no equality symbol, no metric is prescribed, the a priori bounds need not have a common upper bound, and no modulus of uniform continuity is required, in contrast with the metric structures of continuous logic [BYBHU08]; the two-valued counterpart is the logic without equality studied by Keisler and Miller [KM01, Section 4]. General structures admit the lattice operations with their natural interpretation, although these operations are not continuous for the operator-norm distance between projections: the meet of two distinct rays is the zero subspace however close the rays are (Remark 10.2).

Our theory T gives the full description. The theory has sorts for vectors, for events and for square-summable sequences of vectors, and a predicate μ for the probability of an event (Sections 4 and 5). Countable suprema and infima express the projection onto the meet of two events, the closed span of a sequence, the spanning of the space by a named orthonormal sequence, and the countable additivity of μ. The standard models of T[2] represent

[1] A model of our theory T is standard when it is built on full Hilbert spaces, in a sense made precise below; the term is unrelated to the standard models of arithmetic and to nonstandard analysis.

[2] Precisely, a model of T is standard when its vector and sequence sorts are full closed unit balls of complex Hilbert spaces. Nonstandard models exist, among them the countable models given by the downward Löwenheim–Skolem theorem. The vector sort of every model maps into the closed unit ball of a complex Hilbert space by a map that preserves the named affine operations, the inner product and the norm (Lemma 5.3); in

exactly the density-operator measures on the lattice of closed subspaces of a separable infinite-dimensional complex Hilbert space (Theorem B).

The increased expressive power afforded by infinitary sentences comes at a price: compactness is no longer available. The dispersion-free hypothesis, the single closed condition that every probability is 0 or 1, yields together with T a theory T^{01} with no models, standard or not, by a proof that uses neither Gleason's theorem nor standardness (Theorem C). Yet every finite subset of the axioms of T^{01} has a model (Proposition 8.3).

Our refutation theorem shows that the obstructions nevertheless remain finite. A countable theory of closed conditions, each of which asserts that the value of a sentence is at most or at least a given rational number, is inconsistent exactly when a well-founded tree refutes it; the leaves of the tree are finite sets of commitments, placing the values of sentences in closed intervals with rational endpoints, that no structure satisfies together (Theorem A). Completeness theorems are known for deductive calculi with infinitary rules, in two-valued $L_{\omega_1\omega}$ [Kar64, LE65] and in Keisler's probability logic L_{AP} [Hoo78, Kei85]. Our refutation theorem is of a different kind: it concerns real-valued sentences interpreted in general structures, and its certificates are winning strategies in a game, whose plays end in finite contradictions. The trees arise from a topological two-player game of model existence for real-valued logic (Section 3); for two-valued logic, Väänänen's model existence game plays the corresponding role [Vää11]. The players are Olivia and Josh: Olivia maintains that the theory has no model, while Josh attempts to build one, committing the values of sentences to closed intervals with rational endpoints and answering Olivia's challenges about witnesses, function values, and countable suprema and infima. Models are built from Henkin types [ADI26, Lemma 3] along fair plays, in which every challenge that becomes available is eventually made, and the plays that Olivia wins form an open set in the topology generated by finite positions. The theorem of Gale and Stewart [GS53] therefore gives Olivia a winning strategy exactly when the theory is inconsistent, and the plays that follow her strategy form the refuting tree. Only challenges about countable suprema and infima give Josh infinitely many answers, and a well-founded tree that branches only finitely is finite by Kőnig's lemma, in which case some finite set of axioms is already inconsistent. For T^{01}, every refuting tree therefore has an infinitely branching node (Corollary 8.4); for the axioms of T as presented, the finiteness resides in the terminal contradictions, whose commitments concern sentences that may themselves be infinitary. In the Polish space of plays that follow a winning strategy of Josh, the fair plays form a dense G_δ set, and Olivia's bookkeeping is a winning strategy for the second player in the associated Banach–Mazur game (Section B). The category argument is distinct from the one for omitting types, which goes back to Morley [Mor74] and takes place in the Čech-complete space of Henkin types [ADI26]; Section B also relates the game to the Henkin Realizability theorem [ADI26, Proposition 5] and to Keisler's model-theoretic forcing [Kei71, Kei73].

The resulting finiteness theorem is maximal in a precise sense. A hypothesis about hidden variables expressed by countably many closed conditions has no model together with T, standard or not, if and only if a well-founded tree of finite contradictions refutes the two together (Theorem A). The certified hypotheses are therefore exactly the obstructions, and no sound method of refutation can certify more. The criterion concerns all models of T: a hypothesis that fails in every standard model but holds in some nonstandard model is not an obstruction in this sense.

---

a nonstandard model the map may omit vectors of the ball or identify distinct elements, and the sequence sort may fail to be a full Hilbert ball.

The certificates are explicit. Whereas a configuration of Kochen and Specker is chosen in advance and serves against every assignment, in the strategy we construct (Proposition 8.6) the hypothetical measure selects its own configuration. Countable additivity along an orthonormal basis compels Josh to commit some ray to a probability close to 1; the ray determines a four-dimensional subspace, and a parity argument on the eighteen rays of Cabello, Estebaranz and García-Alcaine [CEGA96] closes every branch. The occurrences of a ray in different contexts are distinct terms. Axioms that do not mention μ prove that the terms denote the same projection, and only then does the substitutivity axiom (M3) of Section 5 carry the probability across the identity. The certificate thus records separately the geometric identities and each use of noncontextuality, and its parity contradiction survives explicit numerical tolerances (Theorem 9.2). Whether the least rank of a refuting strategy for T^{01} is finite remains open (Question 8.7).

The earlier logical formulations of contextuality obtain their finite results from compactness. J. L. Bell gave the argument of Kochen and Specker its logical form [Bel96, Theorem 1]. Events become propositional variables, the Boolean operations within a context become connectives, and a noncontextual assignment becomes a satisfying truth assignment. Propositional logic is compact and the constraints on an assignment are closed; Gleason's theorem therefore yields, in every finite dimension at least three, a finite inconsistent set of constraints, and the existence of configurations of Kochen and Specker becomes a consequence of compactness. Propositional logic records only the values 0 and 1, and the probabilities of quantum theory lie outside its language.

Hrushovski and Pitowsky gave a formulation in classical two-valued first-order logic [HP04], in which probabilities are elements of a sort of scalars: their theory has variables for directions and for scalars, and it yields quantitative constraints on quantum probabilities, among them Pitowsky's logical indeterminacy principle [Pit98], together with finite approximations to the predictions of a quantum state. Their formulation rests on the first-order theories of the real and complex fields: compactness yields finite sets of rays on which the probabilities specifying a state constrain further probabilities to any prescribed accuracy, and the decidability of real closed and algebraically closed fields makes the construction effective. The formulation remains within finite dimension, where orthonormal bases are finite.

In the formulations of J. L. Bell and of Hrushovski and Pitowsky, compactness reduces an infinite family of constraints to a finite one, which concerns finitely many rays in finite dimension. Gleason's setting lies beyond propositional and first-order logic: neither expresses the two features of the full setting recalled above, the completeness of $L(H)$ and the countable additivity of the measure. Fixed finite configurations reach the finitely additive part of the subject: infinite-rank projections built from a configuration of Kochen and Specker refute every finitely additive $\{0,1\}$-valued measure on $L(H)$, and Döring extended the Kochen–Specker theorem in this form to von Neumann algebras [Dör05]. For obstructions that rest on countable additivity, no certificate made of finite contradictions had, to our knowledge, been given.

In our setting the refutation theorem takes the place of compactness. The axioms of T express, through countable suprema and infima, the closed span of a sequence of vectors and the countable additivity of the measure, and in the standard models the events form the complete lattice $L(H)$. Every obstruction is refuted, together with T, by a well-founded tree whose terminal contradictions are finite (Theorem A), whereas for T^{01} no finite subset of the axioms is inconsistent and every refuting tree branches infinitely (Proposition 8.3 and Corollary 8.4); the finiteness resides in the terminal contradictions and cannot be transferred to

the axioms. Where Hrushovski and Pitowsky derive finite results from compactness in finite dimension, finite contradictions here certify every obstruction in the infinite-dimensional setting of Gleason's theorem.

We also determine exactly which sets of probabilities the language defines. The measure takes every value in $[0,1]$ in every standard model, and the sets of probabilities uniformly definable across standard models are exactly the σ-projective subsets of $[0,1]$ (Theorem 7.7). The classification of definable sets connects the formalism with a classical chapter of descriptive set theory. The σ-projective sets form the least class that contains the Borel sets and is closed under complements, countable unions and projections, and the class properly contains the projective sets of Lusin and Sierpiński [Lus25, Sie25]. Lusin proved Suslin's analytic sets Lebesgue measurable and observed that the methods did not reach the next projective level [Lus25]. Gödel announced in 1938 that in the constructible universe a projective set of the second level is not Lebesgue measurable [Göd38], whereas under large cardinal hypotheses every σ-projective set is Lebesgue measurable, through the determinacy of the associated games [MS89, AMS21] and the theorem of Mycielski and Świerczkowski [MŚ64]. Our classification is a theorem of ZFC; whether every set of probabilities definable over T is measurable is a question that ZFC does not settle, granted the consistency of the large cardinals involved.

Main results of the paper. We first state the refutation theorem, which concerns every countable theory of closed conditions, and then the theorems on the theory T of quantum probability. The refutation theorem applies to every theory whose axioms are *closed conditions*, that is, assertions $\varphi \leq r$ or $\varphi \geq r$ about sentences φ, with r rational; an axiom with free variables takes this form once its variables are bound by suprema or infima, and the axioms of the theory T described below are all of this kind.

Theorem A (Finite obstructions). *Let T' be a nonempty countable theory of real-valued $L_{\omega_1\omega}$ consisting of closed conditions. Then T' has no models if and only if, in the model existence game of Section 3, the player who claims that T' has no models has a winning strategy. In that case, for every such strategy, the tree of plays in which the same player follows it is well-founded; at each leaf, the commitments made along the play, each of which says that the value of a sentence lies in a closed interval with rational endpoints, form a finite set that is satisfied in no structure. The tree branches infinitely only at the moves that concern countable suprema and infima; once the other player's answers to them are fixed, the plays have bounded length. If the tree branches only finitely, the tree is finite and some finite subset of T' has no models; this is the case when no axiom of T' contains a countable supremum or infimum.*

Each play that follows a winning strategy ends after finitely many rounds, at a finite contradiction, although the theorem supplies no uniform bound on the lengths of the plays. The trees are countable and well-founded, so every winning strategy has a countable ordinal rank, which is finite exactly when the plays that follow the strategy have bounded length; the least rank of a winning strategy measures the complexity of refuting the given theory in the game (Section 3).

The theory T speaks of three kinds of objects, which the logic keeps apart as sorts. The elements of the *Hilbert sort* are vectors. The *Solèr sort* carries the operations of an ortholattice; in the intended interpretation its elements are closed subspaces, each identified, as above, with the orthogonal projection onto it. The elements of the *sequence sort* are sequences

of vectors, through which the completeness of the lattice becomes expressible. A function symbol names the orthogonal projection of a vector onto an element of the Solèr sort, and a vector belongs to such an element when the projection leaves it fixed. The predicate μ assigns a probability to every element of the Solèr sort.

We call a model *standard* when its Hilbert and sequence sorts are full closed unit balls of complex Hilbert spaces with their own operations, inner products and norms (see Definition 4.2 below). The theory T has standard models, and, like every countable theory that has a model, it also has models whose sorts are countable, which are not standard. The Hilbert sort of every model of T is mapped into the closed unit ball of a complex Hilbert space by a map that preserves the named affine operations, the inner product and the norm (see Lemma 5.3). In a nonstandard model, the image of this map may omit vectors of the ball, distinct elements may have the same image, and the sequence sort may fail to be a full Hilbert ball. The term "nonstandard" is here unrelated to nonstandard analysis. Our theorems distinguish the conclusions that hold in every model from those that require standardness.

To state how T captures the setting of Gleason's theorem, consider a standard model, whose Hilbert sort is the closed unit ball of a complex Hilbert space H. For an element a of its Solèr sort, let S_a be the closed subspace of H spanned by the vectors of the Hilbert sort that belong to a, and let P_a be the orthogonal projection onto S_a; for a closed subspace a of H, we likewise write P_a for the orthogonal projection onto a. Equality on the Solèr sort is the relation $[a = b] = 1$ of Section 4, which holds exactly when $S_a = S_b$ (see Lemma 5.4).

Theorem B (Gleason's theorem in real-valued $L_{\omega_1\omega}$). *There is a countable theory T of real-valued $L_{\omega_1\omega}$ with the following two properties.*

(1) In every standard model of T, the Hilbert sort is the closed unit ball of a separable infinite-dimensional complex Hilbert space H, and the map $a \mapsto S_a$ induces an isomorphism of ortholattices from the quotient of the Solèr sort by its equality onto the lattice $L(H)$ of closed subspaces of H; moreover, there is a unique density operator ρ on H such that $\mu(a) = \mathrm{Tr}(\rho P_a)$ for every element a of the Solèr sort.

(2) Conversely, for every separable infinite-dimensional complex Hilbert space H and every density operator ρ on H, there is a standard model of T whose Hilbert sort is the closed unit ball of H, whose Solèr sort is $L(H)$, and in which $\mu(a) = \mathrm{Tr}(\rho P_a)$ for every $a \in L(H)$.

In a separable infinite-dimensional Hilbert space, the probabilities that a quantum state assigns to the closed subspaces fill the interval $[0, 1]$, even when the state is pure, that is, when its density operator is the orthogonal projection onto a ray (see Lemma 7.4 below). Thus the range of μ in every standard model of T is $[0, 1]$ (Proposition 7.5). If a Borel set $B \subseteq \mathbb{R}$ meets $[0, 1]$, adjoining an axiom that excludes B from the range of μ therefore gives a theory with no standard models.

The hypothesis that every probability is 0 or 1 can be written as a single axiom of the logic,

$$\mu(a) \wedge (\mathbf{1} - \mu(a)) \leq 0.$$

In this axiom, the symbol $\mathbf{1}$ denotes the sentence whose value is 1 in every structure, and for formulas φ and ψ the expression $\varphi \wedge \psi$ denotes the minimum of their values. Since the axioms of T confine μ to $[0, 1]$, the axiom says that μ takes only the values 0 and 1. We write

T^{01} for the theory T together with this axiom. Our next theorem shows that the axiom fails not only in the standard models but in every model of T.

Theorem C (The Bell–Kochen–Specker theorem in real-valued $L_{\omega_1\omega}$). *The theory T^{01} has no models, standard or not.*

The proof of Theorem C, given in Section 8, uses neither Gleason's theorem nor the standardness of a model. Every finite set of axioms of T^{01} has a model (Proposition 8.3), so every winning strategy for T^{01} has an infinitely branching node (Corollary 8.4), and no finite set of axioms is a certificate. Whether the least rank of a winning strategy for T^{01} is finite is asked in Question 8.7.

The general theorem applies to every winning strategy; the configuration of rays and the selection of a four-dimensional subspace described above come from the strategy of Section 8.1. In the plays of this strategy that reach the parity argument, the rays are those of the eighteen-ray configuration of Cabello, Estebaranz and García-Alcaine [CEGA96], each of which occurs in exactly two of its nine bases. In the Boolean system that has one Boolean variable for each occurrence of a ray and requires exactly one variable of each basis to be true, the parity contradiction needs all eighteen identities, one between the two occurrences of each ray. No finite set of axioms of T^{01} forces this localization: in the models of finite sets of axioms constructed in the proof of Proposition 8.3, the measure vanishes on every finite-dimensional subspace, as in Proposition 9.1.

In Section 9, after the mathematical development, we give the foundational interpretation. There we distinguish quantum probabilities from the responses of a hidden-variable model, explain why the geometric identities of projection events and the probabilistic requirements on a dispersion-free valuation enter the game through separate axioms, and relate the contextual presentations to earlier approaches.

The full range of the measure also reveals the expressive power of the logic. A subset of $[0,1]$ is *definable over T* when some formula of the logic, with one free variable a of the Solèr sort, singles it out as follows (see Definition 7.6 below). In every standard model of T and at every element a of its Solèr sort, the formula takes the value 1 if $\mu(a)$ lies in the set and the value -1 otherwise. Because formulas may quantify over the Solèr sort, and because the measure takes every value in $[0,1]$, the definable sets form a much larger class than the Borel sets. The σ-projective sets, studied by Aguilera, Müller and Schlicht [AMS21], are the members of the least class of subsets of Polish spaces that contains the Borel sets and is closed under complements, countable unions and projections along $\mathbb{N}^{\mathbb{N}}$. This class contains every projective set, and the example following Theorem 7.7 shows that the containment is strict.

Theorem D (Definable sets of probabilities). *A subset of $[0,1]$ is definable over T if and only if it is σ-projective (see Theorem 7.7).*

Its proof uses both the full range of the measure (Proposition 7.5) and a standard model whose symbols are interpreted by Borel functions on Polish spaces. In that model, a formula of finite quantifier depth k defines a $\mathbf{\Delta}^1_{k+1}$ set; the nonprojective sets therefore require countable joins of formulas of unbounded quantifier depth (Proposition 7.8). The uniformity over standard models is essential: a single formula defines an arbitrary subset of $(0,1)$ in a suitably chosen standard model (Proposition 7.9). By Theorem D, the hypothesis that

the probabilities avoid a given nonempty σ-projective subset of $[0,1]$ can be written as an axiom of the logic, and, by Proposition 7.5, the theory T together with this axiom has no standard models. The exact σ-projective classification, including the strict extension beyond the projective sets, is proved in ZFC.

The classification has a consequence beyond the reach of ZFC. Every σ-projective set belongs to the model $L(\mathbb{R})$, so under large cardinal hypotheses, for instance infinitely many Woodin cardinals with a measurable cardinal above them, every subset of $[0,1]$ definable over T is Lebesgue measurable, has the property of Baire and, if uncountable, contains a perfect set [Lar04, Kec95]. In Gödel's constructible universe, by contrast, some $\mathbf{\Delta}^1_2$ subset of $[0,1]$ is not Lebesgue measurable [Mos09], and this set is definable over T. Whether every set of probabilities that the logic of quantum probability can define is Lebesgue measurable is therefore not decided by ZFC, and large cardinals decide it. The countable joins and the quantifiers over the Solèr sort produce this expressive power together; whether formulas without countable joins can define any set other than $\emptyset$ and $[0,1]$ we do not know (see Remark 10.2).

The principal dependencies among the results are displayed in Figure 1. The representation and definability results concern standard models, whereas the dispersion-free inconsistency holds in every model. The general model existence theorem supplies the passage from this inconsistency to well-founded trees of finite obstructions.

Organization of the paper. In Section 1 we describe the events represented by the theory T and explain, through the obstruction to expressing completeness over the lattice alone (Proposition 1.1), why its vocabulary includes a sort of sequences of vectors. The detailed lattice-theoretic definitions, geometric representation results and their proofs are collected in Section A; their principal sources are MacNeille completions [Mac37], the coordinatization of projective geometries [Bae52, MM70, Hol95], and the theorems on orthomodular spaces due to Amemiya and Araki [AA66] and to Solèr [Sol95]. The logic is presented in Section 2. In Section 3 we organize the model existence theorem as a game, whose analysis rests on the determinacy theorem of Gale and Stewart [GS53], and obtain the well-founded trees of finite obstructions for every countable theory of closed conditions without models, proving Theorem A; that section uses only Section 2. The vocabulary and the axioms of T follow in Sections 4 and 5. The standard models are studied in Sections 6 and 7, which contain the proofs of Theorem B and Proposition 7.5 and of the classification of definable sets of probabilities, Theorem D. In Section 8 we prove Theorem C, that no model of T satisfies the dispersion-free axiom, and follow the proof in the game, where every winning strategy branches infinitely (Corollary 8.4). In Section 9 we compare the finite obstructions with the configurations of Kochen and Specker: what a vertex is, which identifications are made and from which axioms they follow, the tolerances of the intervals, the location of the configuration, the role of σ-additivity for configurations of rays in infinite dimension, and the topological property of the logic behind each kind of certificate. The remarks of Section 10 may be omitted on a first reading. In Section B we relate the game to Keisler's model-theoretic forcing [Kei73] and compare it with the space of Henkin types [ADI26].

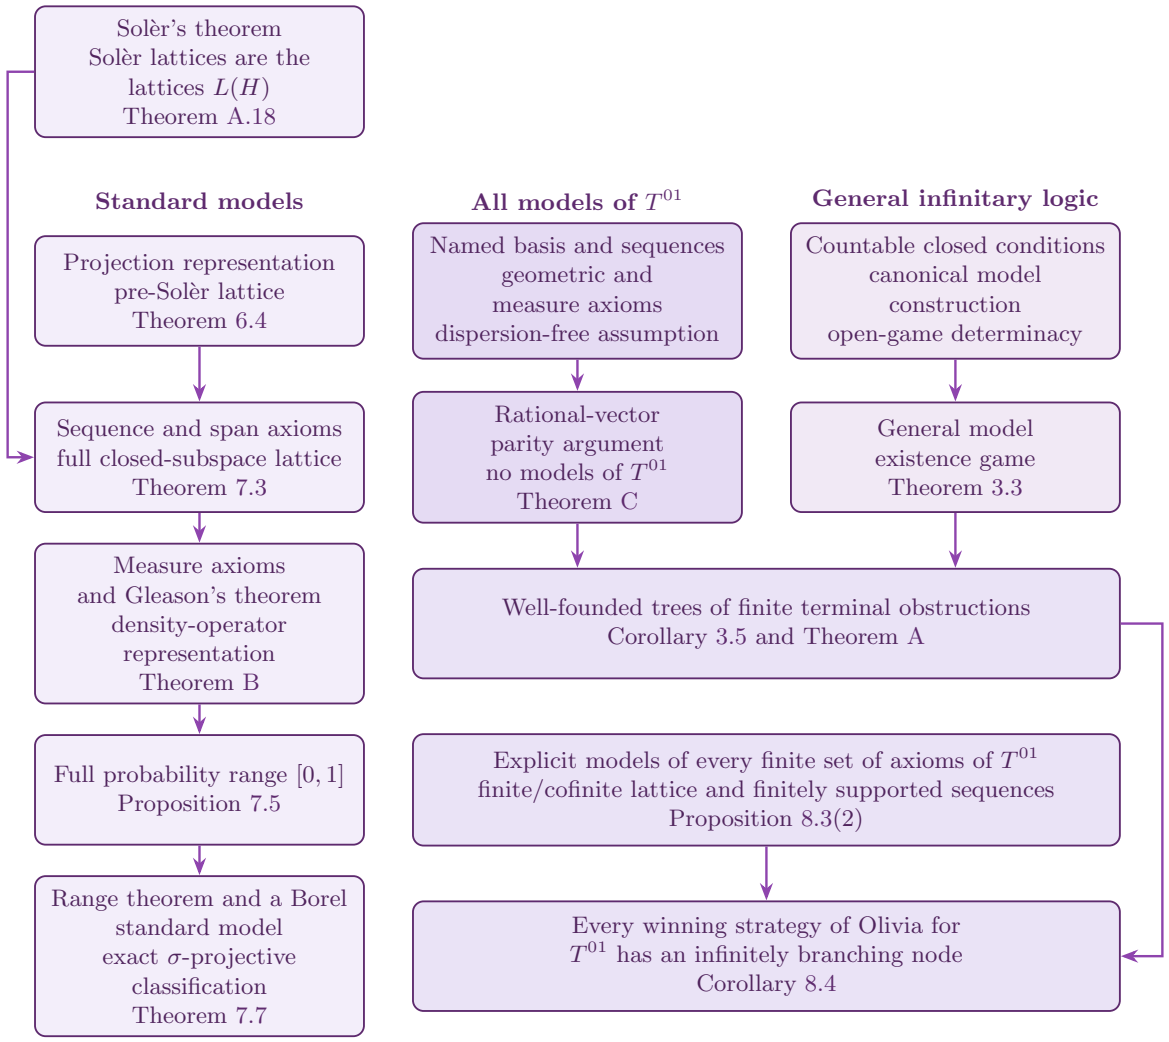


FIGURE 1. Selected direct dependencies among the results; an arrow points from a result to a result that uses it.

1. The geometric setting

The events of a quantum system whose Hilbert space is H are the closed subspaces of H, each determined by the orthogonal projection onto it. Ordered by inclusion, they form the lattice $L(H)$, in which the meet of two subspaces is their intersection, their join is the closed linear span of their union, and the orthocomplement of a subspace is its orthogonal complement. Countable additivity, the hypothesis of Gleason's theorem, concerns countable orthogonal families of events, and a formalization of quantum probability must therefore describe not only the subspaces but also the closed spans of countable families of vectors, on which the hypothesis depends.

An *ortholattice* is a bounded lattice with an order-reversing involution $a \mapsto a^\perp$ such that $a \wedge a^\perp = 0$ and $a \vee a^\perp = 1$. Two elements a and b are *orthogonal* when $a \leq b^\perp$. An *atom* is a nonzero element with no element strictly between it and 0; in $L(H)$ the atoms are the one-dimensional subspaces, or *rays*. A lattice is *complete* when every subset has a join. The lattice $L(H)$ is complete, but a subset of $L(H)$ that is closed under the ortholattice operations can contain every ray without containing every closed subspace. For example, in an infinite-dimensional Hilbert space the finite-dimensional subspaces and their orthogonal complements form such a subset (see Example 6.5).

The theory T keeps vectors, events and sequences of vectors in three sorts. A function symbol describes the projection of a vector onto an event; a vector belongs to an event when this projection leaves it fixed. Since the logic prescribes no equality, two elements of the event sort may represent the same subspace. The geometric axioms define an equality relation on that sort, and the representation statements concern its quotient by this relation. In a standard model of the axioms on vectors and events alone, this quotient is a family of closed subspaces of the given complex Hilbert space, closed under the ortholattice operations and containing every ray, but possibly omitting other closed subspaces (Theorem 6.4). The sequence sort and its axioms supply the missing closed spans: in a standard model of the full theory, the quotient is the whole lattice $L(H)$ (Theorem 7.3). Standardness specifies the interpretation of the vector sorts and is not a consequence of the axioms.

The event quotients just described are pre-Solèr lattices, whose definition is given in Section A. That appendix gives the coordinatization theorem and the theorem of Piron, Amemiya and Araki, and Solèr that characterizes the complete ones as lattices of closed subspaces of infinite-dimensional real, complex or quaternionic Hilbert spaces (Theorem A.18). It also describes their MacNeille completions and the obstruction to orthomodularity of those completions. The distinction between an incomplete event lattice and the full lattice is needed here for a logical reason: completeness cannot be expressed over the lattice alone.

Proposition 1.1. *In the vocabulary of ortholattices expanded by constants p_n, for $n \in \mathbb{N}$, no sentence of classical two-valued $L_{\omega_1\omega}$ is equivalent to completeness among ortholattices in which the p_n are mutually orthogonal nonzero elements.*

Proof. Suppose that a sentence φ expressed completeness in this class. The lattice $L(\ell^2)$, expanded by the rays $p_n = \mathbb{C}e_n$ of its standard orthonormal basis, satisfies φ. Apply the downward Löwenheim–Skolem theorem to a countable fragment containing φ and sentences expressing the ortholattice axioms and the stated conditions on the p_n. It gives a countable ortholattice Λ in the same class that satisfies φ, and hence is complete. For each subset $A \subseteq \mathbb{N}$, completeness supplies $q_A = \bigvee_{n\in A} p_n$. If $n \notin A$, every p_m with $m \in A$ lies below $p_n^\perp$, so $q_A \leq p_n^\perp$; since $p_n \neq 0$, it follows that $p_n \not\leq q_A$. If $n \in A$, then $p_n \leq q_A$. Thus distinct subsets give distinct joins, contradicting the countability of Λ. □

The same obstruction applies in real-valued $L_{\omega_1\omega}$ when the lattice equality and order are represented by two-valued predicates, with their congruence axioms. For a countable fragment that contains these axioms and a sentence proposed to express completeness, the expanded lattice $L(\ell^2)$ has a countable substructure in which every sentence of the fragment keeps its value. The quotient of this substructure by equality is a countable ortholattice with the named orthogonal nonzero elements. Congruence makes every formula invariant under passage to the quotient, so the proposed sentence retains its value there, whereas the quotient cannot be complete by the argument above. A stronger nonexpressibility statement for pre-Solèr lattices, including the condition that their MacNeille completion be orthomodular, is proved in Proposition A.24. The vocabulary of T therefore adds to the lattice a sort of vectors and a sort of sequences of vectors. In standard models, completeness then becomes the requirement that the closed span of every sequence of vectors be an event, which axioms can state; standardness, which no sentence expresses, fixes the vectors.

2. The logic

Probabilities and norms are real quantities, and so are the real and imaginary parts of inner products. For a finitely additive measure, countable additivity amounts to continuity: the measure of the join of a sequence of events is the supremum of the measures of its finite joins. Real-valued $L_{\omega_1\omega}$ expresses these quantities and this continuity, since its formulas take real values and countable suprema are among its connectives. We recall this logic and the Henkin types from which its models are built; the model existence game of Section 3 uses nothing beyond this section.

2.1. **Real-valued $L_{\omega_1\omega}$.** We use real-valued infinitary logic [ADI26], the infinitary extension of Keisler's model theory of general structures [Kei23, Section 2.1]. We recall its definitions as far as we use them.

Vocabularies and structures. A vocabulary consists of sorts, of function symbols, each with the sorts of its arguments and of its values, and of predicate symbols. Each predicate symbol R comes with the sorts of its arguments and with a nonnegative real number r_R, its *a priori bound*. In this paper the a priori bound of every predicate symbol is a positive integer. The vocabulary contains a nullary predicate symbol $\mathbf{1}$, the *logical unity constant*, with a priori bound 1. A structure for the vocabulary interprets each sort as a nonempty set, each function symbol as a function from the product of its argument sorts to its value sort, and each predicate symbol R as a function from the product of its argument sorts to the interval $[-r_R, r_R]$; the logical unity constant is interpreted as 1. No topology is prescribed with respect to which the functions or the predicates must be continuous, and the vocabulary has no equality symbol.

Formulas. The atomic formulas are the expressions $R(\bar{t})$, where R is a predicate symbol and $\bar{t}$ is a tuple of terms of the appropriate sorts. Formulas are built from the atomic formulas by the connectives $s \cdot \varphi$, for $s \in \mathbb{Q}$, $\varphi + \psi$ and $\varphi \vee \psi$, by the connectives $\sup_x$, one for each variable x, and by countable joins $\bigvee_n \varphi_n$. A countable join is formed from a family $(\varphi_n)_{n\in\mathbb{N}}$ of formulas whose free variables lie in one finite set and whose a priori bounds, defined next, have a finite supremum. The a priori bound of $R(\bar{t})$ is r_R; the a priori bounds of $s \cdot \varphi$, $\varphi + \psi$, $\varphi \vee \psi$, $\sup_x \varphi$ and $\bigvee_n \varphi_n$ are $|s| r_\varphi$, $r_\varphi + r_\psi$, $\max(r_\varphi, r_\psi)$, r_φ and $\sup_n r_{\varphi_n}$, respectively. In a structure, the connectives $s \cdot \varphi$, $\varphi + \psi$ and $\varphi \vee \psi$ are interpreted as multiplication by s, as the sum and as the maximum, the connective $\sup_x$ as the supremum over the sort of x, and a countable join as the supremum of the values of its members. By induction on formulas, every formula φ takes its values in $[-r_\varphi, r_\varphi]$ in every structure, and hence in the interval

$I_\varphi = [-\bar{r}_\varphi, \bar{r}_\varphi]$, where the number $\bar{r}_\varphi$ is the least integer not smaller than r_φ. The number r_φ does not change when terms are substituted for the free variables of φ.

Abbreviations. For a rational number q we write q for the sentence $q \cdot \mathbf{1}$, whose value is q in every structure; in particular, we write 0 for $0 \cdot \mathbf{1}$. We write $-\varphi$ for $(-1) \cdot \varphi$ and $\varphi - \psi$ for $\varphi + (-\psi)$. The minimum, the infimum over the sort of x and the countable infimum are expressed by $\varphi \wedge \psi$, which abbreviates $-(-\varphi \vee -\psi)$, by $\inf_x \varphi$, which abbreviates $-\sup_x -\varphi$, and by $\bigwedge_n \varphi_n$, which abbreviates $-\bigvee_n -\varphi_n$. Finally, the absolute value $|\varphi|$ abbreviates $\varphi \vee -\varphi$. The connectives do not include multiplication, but the square of a formula can be written as a countable join of affine expressions in it, and the product of two formulas as a rational combination of two squares. For a formula φ, let φ^2 denote the countable join of the formulas $(2q) \cdot \varphi - q^2$, for the rational numbers q with $|q| \leq r_\varphi$.

Lemma 2.1. *For every formula φ, the join φ^2 is a formula with a priori bound at most $3r_\varphi^2$, and its value in every structure, at every tuple, is the square of the value of φ.*

Proof. The member of index q has a priori bound $2|q|r_\varphi + q^2 \leq 3r_\varphi^2$, and its free variables are those of φ; the join is therefore a formula. The values of φ lie in $[-r_\varphi, r_\varphi]$, and for a real number u in this interval we have $2qu - q^2 = u^2 - (u-q)^2 \leq u^2$; since the rational numbers of the interval are dense in it, the term $(u-q)^2$ can be made arbitrarily small, and the supremum of the values of the members at u is therefore u^2. □

For formulas φ and ψ, the product $\varphi\psi$ abbreviates $\frac{1}{4}(\varphi+\psi)^2 - \frac{1}{4}(\varphi-\psi)^2$, whose value is the product of the values of φ and ψ. Polynomial expressions with rational coefficients in formulas are therefore formulas.

Conditions and theories. The statements of the logic are not formulas but conditions, which compare the value of a sentence with a rational number. For a sentence φ and a rational number r, a *condition* is an expression $\varphi > r$ or $\varphi < r$. A structure $\mathfrak{M}$ satisfies $\varphi > r$ when the value $\varphi^{\mathfrak{M}}$ of φ in $\mathfrak{M}$ is greater than r, and it satisfies $\varphi < r$ when this value is less than r. A *theory* is a set of conditions, and a *model* of a theory is a structure that satisfies every condition in it. Non-strict inequalities are recovered as countable theories: the *closed condition* $\varphi \geq r$ denotes the theory $\{\varphi > r - \frac{1}{n} : n \geq 1\}$, the closed condition $\varphi \leq r$ denotes $\{\varphi < r + \frac{1}{n} : n \geq 1\}$, and the closed condition $\varphi = r$ denotes the union of the two. An *axiom* is a closed condition on a formula $\varphi(\bar{x})$, possibly without free variables, required of all values of $\bar{x}$. The axiom $\varphi(\bar{x}) \leq r$ is the closed condition $\sup_{\bar{x}} \varphi(\bar{x}) \leq r$, and the axiom $\varphi(\bar{x}) \geq r$ is $\inf_{\bar{x}} \varphi(\bar{x}) \geq r$; the axiom $\varphi(\bar{x}) = r$ is the union of the two. For formulas φ and ψ, we write $\varphi \leq \psi$ and $\varphi = \psi$ for the axioms $\varphi - \psi \leq 0$ and $\varphi - \psi = 0$, which are closed conditions when φ and ψ are sentences.

Truth values. As in the work of Aubone, Dueñez and Iovino [ADI26], the value 1 stands for true and the value -1 for false. A formula is *two-valued* in a structure if it takes only the values 1 and -1 there. For formulas θ and θ' that are two-valued in a structure, negation, disjunction and conjunction are expressed there by $-\theta$, $\theta \vee \theta'$ and $\theta \wedge \theta'$. The existential and the universal quantifier are expressed by $\sup_x$ and $\inf_x$, and countable disjunction and conjunction, for families whose a priori bounds have a finite supremum, by $\bigvee_n$ and $\bigwedge_n$. A supremum of values in $\{-1, 1\}$ over a nonempty set is attained, and the connectives $\sup_x$ and $\inf_x$, applied to a two-valued formula, are therefore the traditional quantifiers over the sort of x. We use them in this way over the Solèr sort, on which equality and order are given by formulas that are two-valued in every model of the theory T_0 of Section 5 (see Lemma 5.4 below). A two-valued formula θ is asserted by the axiom $\theta \geq 1$, and the implication from θ to θ' is asserted by the axiom $-\theta \vee \theta' \geq 1$. If θ is two-valued and ψ is any formula, the axiom

$\theta \wedge \psi \leq 0$ says that θ implies $\psi \leq 0$: where θ is false, the minimum is at most -1. Norms, the measure and the real and imaginary parts of inner products keep their numerical values; every statement about these quantities is a closed condition or an axiom with a rational threshold. A measure or an assignment that takes only the values 0 and 1 is called $\{0,1\}$-valued, never two-valued, since the latter term refers to the truth values 1 and -1.

Truncations. A countable join requires a common a priori bound on its member formulas, not merely a common bound on their values. For example, the formulas $((n \cdot \varphi) \wedge 1) \vee 0$ take values in $[0,1]$, but their a priori bounds are $\max(nr_\varphi, 1)$. When $r_\varphi > 0$, these bounds grow without limit, so the formulas cannot be joined over n.

For the atomic predicates used below, we therefore name their truncations separately in the vocabulary of Section 4, each with a priori bound 1. The axioms (Tr) of Section 5 identify each symbol with the corresponding truncated expression, so the countable joins of the symbols are legitimate formulas that give exact tests for strict inequalities in structures satisfying (Tr).[3]

Following Aubone, Dueñez and Iovino [ADI26], we call a *fragment* a set of formulas that contains the first-order formulas, those in which all joins are finite, and is closed under subformulas, under substitution of terms and under the connectives $s \cdot \varphi$, $\varphi + \psi$, $\varphi \vee \psi$ and $\sup_x$.

2.2. **Henkin types and model existence.** In the model existence argument of Section 3, we build a model from an assignment of values to sentences that respects the semantics of the logic. We recall the definition and the Henkin conditions from the topological treatment of model existence by Aubone, Dueñez and Iovino [ADI26]: an assignment satisfying these conditions is a Henkin type and comes from a structure whose elements are named by closed terms. We record these conditions for the later game and its comparison with the topological construction in Section B; the geometric representation arguments that follow do not use them.

For a countable fragment $\mathcal{F}$ of real-valued $L_{\omega_1\omega}$ for a countable vocabulary, such as the vocabulary $\mathcal{L}$ of Section 4, let C be a set of new constants, countably many of each sort. Let $\mathcal{F}_C$ be the fragment that $\mathcal{F}$ generates in the vocabulary enlarged by C, and let $S_C^{\mathcal{F}}$ be the set of sentences of $\mathcal{F}_C$. An $S_C^{\mathcal{F}}$-*Henkin type* is an element $\mathfrak{m}$ of the product of the intervals $[-r_\varphi, r_\varphi]$, over the sentences φ of $S_C^{\mathcal{F}}$, that satisfies the following *Henkin conditions*:

(H1′) $\mathbf{1}^{\mathfrak{m}} = 1$;
(H2′) $(s \cdot \varphi)^{\mathfrak{m}} = s\,\varphi^{\mathfrak{m}}$ for every rational number s;
(H3′) $(\varphi + \psi)^{\mathfrak{m}} = \varphi^{\mathfrak{m}} + \psi^{\mathfrak{m}}$;
(H4′) $(\varphi \vee \psi)^{\mathfrak{m}} = \max(\varphi^{\mathfrak{m}}, \psi^{\mathfrak{m}})$;
(H5′) $(\sup_x \varphi)^{\mathfrak{m}} = \sup_t \varphi(t)^{\mathfrak{m}}$, the supremum being taken over the closed terms t of the sort of x;
(H6′) $(\bigvee_n \varphi_n)^{\mathfrak{m}} = \sup_n \varphi_n^{\mathfrak{m}}$.

Here $\varphi^{\mathfrak{m}}$ denotes the coordinate of $\mathfrak{m}$ at φ, and the sentences involved range over $S_C^{\mathcal{F}}$. With the topology induced by the product topology, the $S_C^{\mathcal{F}}$-Henkin types form the *space of Henkin*

[3] A different presentation of a normed space can use sorts for all closed balls of positive integer radii, with inclusion maps between the balls and operations between the appropriate sorts. On a ball of radius n the norm has bound n, and on two Hilbert balls of radii n and m the real and imaginary parts of the inner product have bound nm. A scaling map from the unit ball to a larger ball, followed by a separately named truncated predicate of bound 1 there, can supply a steep predicate on the unit ball with bound 1. Here we name the truncations directly; for the measure on the Solèr sort they are likewise new predicate symbols whose values (Tr) fixes.

types $\mathcal{H}^{\mathcal{F}}_C$, a G_δ subset of the product and therefore a Čech-complete space, hence a Baire space by Čech's form of the Baire category theorem [ADI26, Proposition 6 and Corollary 7]. Every Henkin type is the $S^{\mathcal{F}}_C$-type of a structure for the enlarged vocabulary in which every element is the value of a closed term [ADI26, Lemma 3].

Historical note. Carol Karp, a student of Henkin's, wrote her dissertation of 1959 on languages with expressions of infinite length and published a book on them in 1964 [Kar64]. In 1969 Makkai applied Smullyan's method of consistency properties to $L_{\omega_1\omega}$, building models in the manner of Henkin's completeness proof, and extended it to the infinitary logics on admissible sets [Mak69b, Mak69a]. In 1971 Keisler made the model existence theorem for consistency properties a central method of his book on $L_{\omega_1\omega}$ [Kei71, Chapter 3], which also contains the omitting types theorem for countable fragments [Kei71, Chapter 11]. In 1973 Keisler built countable models by forcing with finite sets of sentences satisfiable in a given class of models [Kei73, p. 100, Example 1.2]. In the same work he replaced principal types by $\forall\exists$ formulas over a set of formulas of a countable fragment [Kei73, p. 107]. His general omitting types theorem gives conditions under which a $\forall\exists$ class contains a countable model in which countably many given $\forall\exists$ sentences hold [Kei73, p. 108]. Keisler's finite conditions and $\forall\exists$ formulas have counterparts in the model existence game of Section 3 and in the contextual presentations of Section 9 (see Section B). In 1974 Morley derived the omitting types theorem for countable fragments from the Baire category theorem, applied to a space of countable structures, which he showed to be a Baire space [Mor74]. In 2001 Keisler and Miller carried the theorem into first-order logic without equality, where its usual formulation does not hold. There the model is built from closed terms, and principal types give way to locally realized types [KM01, Section 4].

Model theory with truth values in the unit interval goes back to Chang's work of 1958 and 1961 on the infinite-valued logic of Łukasiewicz [Cha58, Cha61]. The continuous model theory of Chang and Keisler of 1966 took its truth values in a compact Hausdorff space [CK66]. Henson introduced the logic of positive bounded formulas in 1976 to study a Banach space through its nonstandard hull [Hen76]. In metric approaches to real-valued model theory, a distance predicate takes the place of equality, and continuity requirements ensure that the other symbols respect the resulting uniform structure.[4] Uniform continuity with moduli prescribed by the signature was the setting of Iovino's thesis of 1994 and of his papers of 1999 on stable Banach spaces [Iov94, Iov99a, Iov99b]; in 2002 Henson and Iovino presented the model theory of normed space structures in it [HI02, Sections 2 and 3]. In 2008 Ben Yaacov, Berenstein, Henson and Usvyatsov set out the model theory of metric structures in continuous logic [BYBHU08]. In 2009 Ben Yaacov and Iovino introduced an infinitary continuous logic for model-theoretic forcing [BYI09]. In 2010 Ben Yaacov and Usvyatsov built moduli of uniform continuity into the signatures of continuous logic [BYU10, Section 2]. In 2014 Eagle proved the omitting types theorem for infinitary $[0,1]$-valued logic, whose metric structures have a master metric with respect to which the predicates are uniformly continuous [Eag14, Definitions 1.4 and 1.5]. In the model theory for real-valued structures that Keisler set out in 2023, a general structure requires no metric signature: its function and predicate symbols may be interpreted by arbitrary mappings, subject only to the boundedness of the predicates [Kei23, Sections 1 and 2.1]. General structures correspond to first-order structures without equality, as metric structures with a distance predicate correspond to first-order

[4] In Łukasiewicz predicate logic, the congruence axioms for a similarity, or graded equality, make every predicate 1-Lipschitz for the distance given by one minus the similarity (see [Háj98, Section 5.6] and the earlier work collected there).

structures with equality [Kei23, Section 1]. For vocabularies with countably many function and predicate symbols, every general structure expands to one with a distance predicate that is a uniform limit of formulas [Kei23, Theorem 3.3.4]. The logic of this paper, whose conditions are strict inequalities, is the infinitary extension of Keisler's model theory for real-valued structures; Aubone, Dueñez and Iovino introduced it, built its models from closed terms and proved its omitting types theorem by Morley's method in the Čech-complete space of Henkin types [ADI26].

3. The model existence game and finite obstructions

In first-order logic, compactness turns a theory without models into a finite set of axioms without models, much as finitely many vectors suffice in the Kochen–Specker theorem [KS67]. The countable suprema and infima of real-valued $L_{\omega_1\omega}$, through which σ-additivity is expressed, carry the logic beyond the reach of compactness, and in this section we prove, for every countable theory of closed conditions, the theorem that takes its place, Theorem A of the Introduction. To obtain finite obstructions in this setting, we organize the model existence argument of Aubone, Dueñez and Iovino [ADI26, Lemma 3] as a game in which each play of a winning strategy reaches a finite contradiction. The argument realizes Henkin types, the value assignments of Section 2.2, through the Henkin construction of canonical structures whose elements are named by closed terms. The same construction underlies the omitting types theorem of Aubone, Dueñez and Iovino, whose Baire category proof we compare with the game in Section B.

The two players are Olivia, who claims that a given countable theory has no models, and Josh, who attempts to describe one by answering her challenges. Each answer adds information that all later answers must retain (see Figure 2). Olivia wins when Josh has no legal answer or the accumulated intervals for a basic predicate value have no common point in its prescribed range; Josh wins if he can continue indefinitely. The game follows Väänänen's model existence game for two-valued logic [Vää11]. We compare its positions and rules with Keisler's forcing in Section B.

In Theorem 3.3 below we show that a nonempty countable theory of closed conditions has no models exactly when Olivia has a winning strategy. The plays that follow such a strategy form a well-founded tree, with a finite obstruction at each leaf: statements about sentence values that no structure can satisfy together. Only challenges to countable suprema or infima can give Josh infinitely many answers; for a fixed winning strategy, the plays whose countable answers form any one prescribed finite sequence have bounded length. In Section 8.1 we return to the quantum theory: for the dispersion-free extension of the theory T of Theorem B, every finite set of axioms has a model (Proposition 8.3), and in the strategy that follows the proof of Theorem C, Olivia first locates an atom of large measure and then records enough of the finite geometry to run the parity argument. Plays that reach this argument end at positions whose commitments concern a Kochen–Specker configuration; the other plays end at contradictions exposed along the way.

3.1. **The setting.** In this section we work with a countable many-sorted vocabulary $\mathcal{L}$ for the logic of Section 2 and with a nonempty countable theory T in $\mathcal{L}$; within the section, these letters denote an arbitrary vocabulary and theory, not the vocabulary of Section 4 and the theory of Theorem B. As usual, every sort of a structure is nonempty.

In the game, Josh commits the values of sentences to intervals, and the rules require for each sentence an interval fixed in advance that contains every possible value of the sentence. These are the intervals $I_\varphi = [-\bar{r}_\varphi, \bar{r}_\varphi]$ of Section 2, whose integer endpoints are determined

by the a priori bound r_φ and do not change when terms are substituted for the free variables of φ. In the game, the infima $\inf_x$ over a sort and the countable infima $\bigwedge_n$, which abbreviate $-\sup_x -\varphi$ and $-\bigvee_n -\varphi_n$, are treated as primitive, with rules of their own, and the interval of $\inf_x \varphi$ or $\bigwedge_n \varphi_n$ is that of $\sup_x \varphi$ or $\bigvee_n \varphi_n$. The formulas $s\cdot\varphi$, $\varphi+\psi$ and $\varphi\vee\psi$, by contrast, are handled by the rule (Cn) of Definition 3.2 through the functions $t \mapsto st$, $(t,u) \mapsto t+u$ and $(t,u) \mapsto \max(t,u)$ that interpret them. We call these functions the *connective functions*; each of them is continuous. Finally, the logical unity constant $\mathbf{1}$, a nullary predicate symbol, takes only the value 1. For the definition of infeasible positions below, we write $\mathcal{V}_{\mathbf{1}} = \{1\}$ and $\mathcal{V}_R = [-r_R, r_R]$ for every other predicate symbol R.

The axioms enter the game through one of its rules, which is simplest to state when all axioms have the same form. As everywhere in the paper, an axiom of T is a closed condition required of all values of its free variables. If $\bar{x}$ lists the free variables of a formula φ, then, by the conventions of Section 2, an axiom $\varphi(\bar{x}) \le r$ is the closed condition $\sup_{\bar{x}} \varphi \le r$ and an axiom $\varphi(\bar{x}) \ge r$ is $\inf_{\bar{x}} \varphi \ge r$. An axiom $\varphi(\bar{x}) = r$ is the union of these two closed conditions. Every axiom of T may therefore be regarded as a closed condition $\psi \le r$ or $\psi \ge r$ for a sentence ψ and a rational number r, an axiom $\varphi(\bar{x}) = r$ counting as two. The one or two closed conditions on sentences that thus correspond to an axiom constitute its *normal form*. Theories that contain strict conditions are discussed in Remark 3.6.

Josh's description needs names for elements before their predicate values and function images can be tested. For each sort $\mathbf{S}$, we therefore add to $\mathcal{L}$ a set $\mathcal{W}_{\mathbf{S}} = \{w_{\mathbf{S},i} : i \in \mathbb{N}\}$ of new constants of sort $\mathbf{S}$, the *witnesses*. We write $\mathcal{W}$ for their union, $\mathcal{L}_{\mathcal{W}}$ for the expanded vocabulary, and w, w' for witnesses. The game treats a constant symbol of $\mathcal{L}$ as a function symbol with no arguments, so that its value is named by the same rule as any other function output. A sentence of $\mathcal{L}_{\mathcal{W}}$ is *basic* when it has the form $R(\bar{w})$, where R is a predicate symbol and $\bar{w}$ is a tuple of witnesses of the appropriate sorts. The tuple may be empty; in particular, the sentence $\mathbf{1}$ is basic.

The description retains both numerical bounds and the function values needed to interpret them. A commitment places a sentence value in an interval, while a record names the output of a function at given witnesses. The axiom rule introduces commitments, and the totality rule supplies function records. The other rules test commitments by following the definition of sentence values: a connective leads to its arguments, a countable supremum or infimum to its members, and a supremum or infimum over a sort to instances at witnesses. For an atomic sentence, the records replace compound terms by witnesses until a basic sentence remains. Thus a conflict among composite assertions can eventually be exposed as incompatible intervals for a basic predicate value.

Definition 3.1. A *commitment* is an expression $\varphi \in I$, where φ is a sentence of $\mathcal{L}_{\mathcal{W}}$ in which only finitely many witnesses occur and $I = [\alpha, \beta]$ is an interval with rational endpoints $\alpha \le \beta$; the interval may be a single point. A *record* is an expression $F(\bar{w}) \mapsto w'$, where F is a function symbol and $\bar{w}, w'$ are witnesses of the appropriate sorts. A *position* is a finite set $\mathfrak{p}$ of commitments and records that contains at most one record $F(\bar{w}) \mapsto w'$ for each function symbol F and each tuple $\bar{w}$. A witness *occurs* in $\mathfrak{p}$ when it occurs in a commitment or a record of $\mathfrak{p}$. In a structure for $\mathcal{L}$, a position is *satisfied* under an interpretation of the witnesses that occur in it when each of its commitments holds and, for each of its records $F(\bar{w}) \mapsto w'$, the term $F(\bar{w})$ and the witness w' have the same value. A position $\mathfrak{p}$ is *infeasible* when there is a basic sentence $R(\bar{w})$ such that the intervals I for which $\mathfrak{p}$ contains the commitment $R(\bar{w}) \in I$, together with $\mathcal{V}_R$, have empty intersection.

An infeasible position is contradictory in the most elementary way: it assigns to a basic sentence $R(\bar{w})$ finitely many rational intervals that have no common point in $\mathcal{V}_R$ (see Figure 4). Infeasibility is the elementary conflict detected by the rules; satisfiability asks whether all the commitments and records can hold together in a structure. A position may already be unsatisfiable without being infeasible, and Olivia's subsequent challenges may expose the contradiction at the level of basic sentences.

To test a commitment on a connective, Olivia needs bounds on its arguments that control the resulting value. Josh therefore answers with small grid cells rather than exact real numbers. The continuity of the connective on the product of the fixed intervals assigned to its argument sentences determines how fine the cells must be; a prescribed widening of the committed interval allows for their width. For $m \in \mathbb{N}$, the *dyadic grid* $\mathcal{G}_m$ is the set of intervals $[i2^{-m}, (i+1)2^{-m}]$, for $i \in \mathbb{Z}$. Let ℓ be 1 or 2, let $f\colon \mathbb{R}^\ell \to \mathbb{R}$ be a connective function, let $\varphi_1, \dots, \varphi_\ell$ be sentences, and let $k \in \mathbb{N}$ be given. The function f is uniformly continuous on the compact box $I_{\varphi_1} \times \cdots \times I_{\varphi_\ell}$; we write $m(f, \bar{\varphi}, k)$ for the least $m \in \mathbb{N}$ such that $|f(\xi) - f(\xi')| \le 2^{-k}$ for all points ξ, ξ' of the box with $|\xi_j - \xi'_j| \le 2^{-m}$ for $j = 1, \dots, \ell$. For an interval $I = [\alpha, \beta]$ and $k \in \mathbb{N}$, we write $I^{(k)} = [\alpha - 2^{-k}, \beta + 2^{-k}]$.

When Josh chooses a witness, the rules allow witnesses already in the position and one designated fresh witness. This restriction leaves finitely many choices at each such move, while permitting a previously unused witness whenever one is needed. The choices of an index for a countable supremum or infimum retain their countable branching; this distinction is the basis of the finite-obstruction analysis in the third subsection.

Definition 3.2 (The model existence game). The game $\mathrm{G}(T)$ is played by Olivia and Josh. It starts at the empty position; at each round, Olivia makes a *challenge*, and Josh *answers* it by adding to the current position $\mathfrak{p}$ commitments and records, if any, in the manner that the rule of the challenge prescribes. When Josh answers a challenge, the *first fresh witness* of a sort $\mathbf{S}$ is the witness $w_{\mathbf{S},i}$ with the least index i that occurs neither in $\mathfrak{p}$ nor in the challenge. In each rule other than (Ax) and (Tot), the commitment named by Olivia belongs to $\mathfrak{p}$, and its interval is written $I = [\alpha, \beta]$.

(Ax): Olivia names an axiom $\psi \le r$ or $\psi \ge r$ of T. Josh adds $\psi \in [\min(-\bar{r}_\psi, r), r]$ or $\psi \in [r, \max(r, \bar{r}_\psi)]$, respectively.

(Tot): Olivia names a function symbol F and witnesses $\bar{w}$ of the appropriate sorts. If $\mathfrak{p}$ contains no record $F(\bar{w}) \mapsto w'$, Josh adds such a record, choosing for w' a witness of the sort of the values of F that occurs in $\mathfrak{p}$ or is the first fresh witness of this sort; otherwise he adds nothing.

(Tm): Olivia names a commitment $R(\bar{t}) \in I$, where R is a predicate symbol and $\bar{t}$ is a tuple of terms, together with a term $F(\bar{w})$ that occurs in $\bar{t}$ and in which $\bar{w}$ is a tuple of witnesses (possibly empty). Josh first answers as in (Tot) for F and $\bar{w}$; then, with w' the witness such that the enlarged position contains the record $F(\bar{w}) \mapsto w'$, he adds $R(\bar{t}') \in I$, where $\bar{t}'$ is obtained from $\bar{t}$ by replacing every occurrence of $F(\bar{w})$ by w'.

(Cn): Olivia names a commitment $f(\varphi_1, \dots, \varphi_\ell) \in I$, where f is a connective function, and a number $k \in \mathbb{N}$. With $m = m(f, \bar{\varphi}, k)$, Josh chooses intervals $D_1, \dots, D_\ell$ of $\mathcal{G}_m$ such that each D_j meets I_{φ_j} and f maps $(D_1 \cap I_{\varphi_1}) \times \cdots \times (D_\ell \cap I_{\varphi_\ell})$ into $I^{(k)}$, and adds the commitments $\varphi_j \in D_j \cap I_{\varphi_j}$ (see Figure 3).

($\sup \le$): Olivia names a commitment $\sup_x \varphi(x) \in I$, where x is a variable of sort $\mathbf{S}$, and a witness $w \in \mathcal{W}_{\mathbf{S}}$. Josh adds $\varphi(w) \in [-\bar{r}_\varphi, \beta]$.

(sup ≥): Olivia names a commitment $\sup_x \varphi(x) \in I$, where x is a variable of sort $\mathbf{S}$, and a number $k \in \mathbb{N}$. Josh chooses a witness $w \in \mathcal{W}_{\mathbf{S}}$ that occurs in $\mathfrak{p}$ or is the first fresh witness of sort $\mathbf{S}$, and adds $\varphi(w) \in [\alpha - 2^{-k}, \bar{r}_\varphi]$.

(inf ≥): Olivia names a commitment $\inf_x \varphi(x) \in I$, where x is a variable of sort $\mathbf{S}$, and a witness $w \in \mathcal{W}_{\mathbf{S}}$. Josh adds $\varphi(w) \in [\alpha, \bar{r}_\varphi]$.

(inf ≤): Olivia names a commitment $\inf_x \varphi(x) \in I$, where x is a variable of sort $\mathbf{S}$, and a number $k \in \mathbb{N}$. Josh chooses a witness w as in (sup ≥), and adds $\varphi(w) \in [-\bar{r}_\varphi, \beta + 2^{-k}]$.

($\bigvee \leq$): Olivia names a commitment $\bigvee_n \varphi_n \in I$, where $(\varphi_n)_{n\in\mathbb{N}}$ is a family of sentences, and a number $n \in \mathbb{N}$. Josh adds $\varphi_n \in [-\bar{r}_{\bigvee_n \varphi_n}, \beta]$.

($\bigvee \geq$): Olivia names a commitment $\bigvee_n \varphi_n \in I$ as in the preceding rule, and a number $k \in \mathbb{N}$. Josh chooses a number $n \in \mathbb{N}$ and adds $\varphi_n \in [\alpha - 2^{-k}, \bar{r}_{\bigvee_n \varphi_n}]$.

($\bigwedge \geq$): Olivia names a commitment $\bigwedge_n \varphi_n \in I$, where $(\varphi_n)_{n\in\mathbb{N}}$ is a family of sentences, and a number $n \in \mathbb{N}$. Josh adds $\varphi_n \in [\alpha, \bar{r}_{\bigwedge_n \varphi_n}]$.

($\bigwedge \leq$): Olivia names a commitment $\bigwedge_n \varphi_n \in I$ as in the preceding rule, and a number $k \in \mathbb{N}$. Josh chooses a number $n \in \mathbb{N}$ and adds $\varphi_n \in [-\bar{r}_{\bigwedge_n \varphi_n}, \beta + 2^{-k}]$.

An answer is *legal* when it obeys the rule and every interval that it adds is nonempty, that is, has its left endpoint at most its right endpoint. Josh *loses* at a round if he has no legal answer to Olivia's challenge, or if his answer makes the position infeasible. A *play* is a finite or infinite sequence of rounds played according to these rules; it ends at the round, if any, at which Josh loses, a round that consists of Olivia's challenge alone when he has no legal answer to it. Olivia wins a play if Josh loses at some round; Josh wins a play if he never loses, in which case the play is infinite. A *strategy* for a player chooses the player's next move from the moves made so far; it is *winning* when the player wins every play in which the strategy is followed, whatever the opponent does.

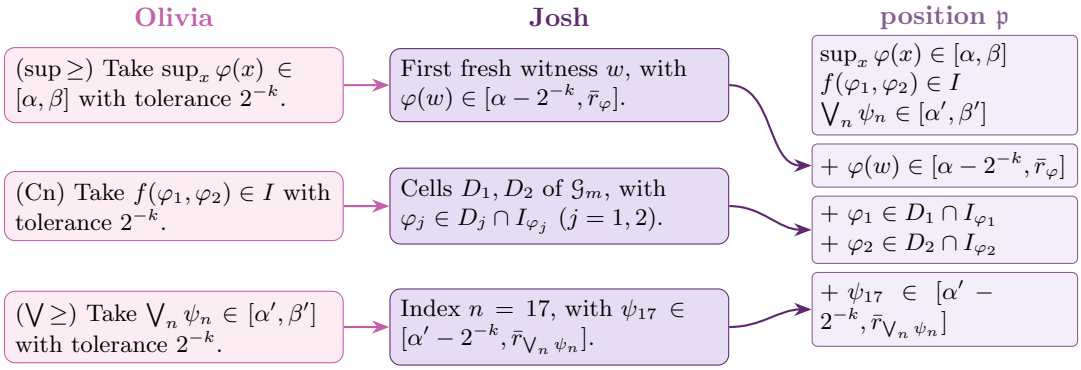


FIGURE 2. Three rounds of the game, the first concerning a supremum over a sort, the second a connective function and the third a countable supremum, with the position growing by one card at each answer.

Josh's answers are of three kinds, according to the number of choices that the rules leave him. The rules (Ax), (sup ≤), (inf ≥), ($\bigvee \leq$) and ($\bigwedge \geq$) leave him no choice: the answer is determined by the challenge. Under (Tot), (Tm), (Cn), (sup ≥) and (inf ≤) he has finitely many choices because only finitely many witnesses occur in a position and only finitely many intervals of $\mathcal{G}_m$ meet a bounded interval. Only the rules ($\bigvee \geq$) and ($\bigwedge \leq$), which concern countable suprema and infima, leave him a choice among infinitely many answers. We call

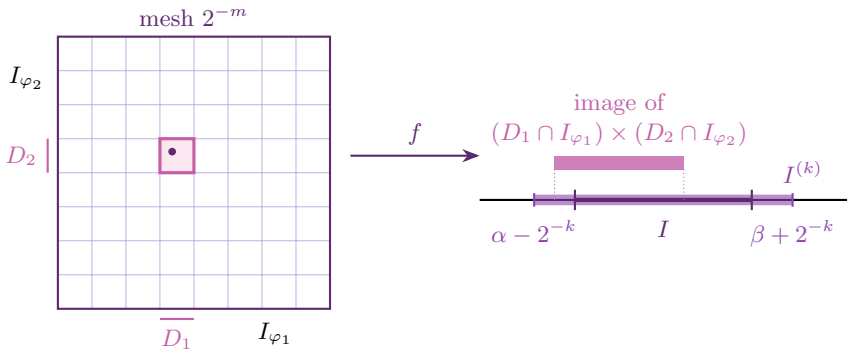


FIGURE 3. The rule (Cn) for a connective f of two variables: Josh must choose a cell $D_1 \times D_2$ of the dyadic grid of mesh 2^{-m} whose intersection with the box $I_{\varphi_1} \times I_{\varphi_2}$ is mapped by f into the widened interval $I^{(k)}$.

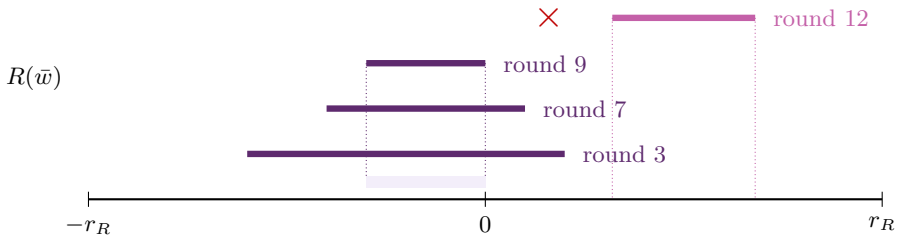


FIGURE 4. An infeasible position: at different rounds, Josh has committed the basic sentence $R(\bar{w})$ to intervals that have no common point.

the challenges made under these two rules *countable challenges*, and Josh's answers to them *countable answers*.

3.2. **The model existence theorem as a game.** Josh's survival yields a model only if every requirement is eventually tested: an axiom left unnamed or a commitment left unchallenged need not hold in the structure built from the play. A challenge is *available* at a round when the rules allow Olivia to make it there, equivalently when the commitment it names, if any, is already in the position. Since positions only grow, an available challenge remains available. The vocabulary $\mathcal{L}_{\mathcal{W}}$ and the theory T are countable, and every sentence occurring in a play is atomic or is obtained from a subformula of an axiom by substituting closed terms for its free variables. Thus only countably many challenges can occur, and at least one is always available because T is nonempty. A play is *fair* if every challenge that becomes available is eventually made. A strategy of Olivia is *fair* if every infinite play in which she follows it is fair, whatever Josh does. The following *bookkeeping* gives a fair strategy. She fixes a bijection $i \mapsto (j(i), l(i))$ from $\mathbb{N}$ onto $\mathbb{N} \times \mathbb{N}$ with $j(i) \leq i$ for every $i \in \mathbb{N}$, and at each round she enumerates, with repetitions if necessary, the challenges then available. At round i she makes the $l(i)$-th challenge of her enumeration for round $j(i)$, which is still available because $j(i) \leq i$.

Theorem 3.3 (The model existence theorem as a game). *Let T be a nonempty countable theory of real-valued $L_{\omega_1\omega}$ whose axioms are closed conditions $\psi \leq r$ or $\psi \geq r$ for sentences ψ and rational numbers r, and let $\mathrm{G}(T)$ be the game of Definition 3.2. The following conditions are equivalent:*

(1) the theory T has a model;
(2) Josh has a winning strategy in $\mathrm{G}(T)$;
(3) Olivia has no winning strategy in $\mathrm{G}(T)$.

Proof. (1) implies (2). Let $\mathfrak{N}$ be a model of T; Josh answers every challenge truthfully in $\mathfrak{N}$. To do so, he keeps along the play a map h from the witnesses occurring in the current position to $\mathfrak{N}$. The map h makes every commitment of the position hold in $\mathfrak{N}$ when each witness w is interpreted as $h(w)$. Moreover, it satisfies the equation $F^{\mathfrak{N}}(h(\bar{w})) = h(w')$ for every record $F(\bar{w}) \mapsto w'$ of the position. When a challenge names a witness that does not yet occur, Josh first extends h to it arbitrarily, which is possible because the sorts of $\mathfrak{N}$ are nonempty. The commitment that he adds under (Ax) holds because $\mathfrak{N}$ is a model of T. Under (Tot) and (Tm), if the position contains no record $F(\bar{w}) \mapsto w'$, Josh adds one with w' the first fresh witness and extends h to w' by letting $h(w') = F^{\mathfrak{N}}(h(\bar{w}))$. Under (Tm), whether the record is new or not, the value of $R(\bar{t}')$ in $\mathfrak{N}$ equals that of $R(\bar{t})$. To a challenge (Cn), Josh answers with cells $D_j \in \mathcal{G}_m$ that contain the values ξ_j of the sentences φ_j in $\mathfrak{N}$. Every point of $(D_1 \cap I_{\varphi_1}) \times \cdots \times (D_\ell \cap I_{\varphi_\ell})$ then has coordinates within 2^{-m} of those of $(\xi_1, \ldots, \xi_\ell)$. By the choice of m, the connective f maps every such point to within 2^{-k} of $f(\xi_1, \ldots, \xi_\ell)$, which lies in I. The commitment added under $(\sup \leq)$ holds because the value of $\varphi(w)$ is at most the supremum, hence at most β. Under $(\sup \geq)$, Josh takes the first fresh witness w. Since the supremum is at least α, he lets $h(w)$ be an element of $\mathfrak{N}$ at which φ takes a value at least $\alpha - 2^{-k}$. The remaining rules are handled in the same way; under $(\bigvee \geq)$, for instance, Josh names a number n for which the value of φ_n is at least $\alpha - 2^{-k}$. Every answer is legal, since the true values lie in the added intervals. Moreover, no position becomes infeasible, since the value in $\mathfrak{N}$ of a basic sentence $R(\bar{w})$ lies in every interval committed for the sentence and in $\mathcal{V}_R$. Since both of these facts hold at every round, Josh never loses.

(2) implies (1). Let Josh play a winning strategy against a fair strategy of Olivia. The resulting play is infinite and fair; let $\mathfrak{p}_0 \subseteq \mathfrak{p}_1 \subseteq \cdots$ be its positions, indexed by $i \in \mathbb{N}$, and let $\mathfrak{p}_\infty$ be their union. We define a structure $\mathfrak{M}$ for $\mathcal{L}_{\mathcal{W}}$ whose elements are the witnesses: the universe of sort $\mathbf{S}$ is $\mathcal{W}_{\mathbf{S}}$, and each witness is interpreted as itself. For a function symbol F and witnesses $\bar{w}$, Olivia makes the challenge (Tot) for F and $\bar{w}$ at some round, because the play is fair. The union $\mathfrak{p}_\infty$ therefore contains a record $F(\bar{w}) \mapsto w'$, which is unique because every $\mathfrak{p}_i$ is a position, and we define $F^{\mathfrak{M}}(\bar{w})$ to be w'. For a basic sentence $R(\bar{w})$, the intervals I for which $\mathfrak{p}_\infty$ contains $R(\bar{w}) \in I$ are compact, and so is $\mathcal{V}_R$. Every finite subfamily of the family formed by these intervals and $\mathcal{V}_R$ has a common point, because no position $\mathfrak{p}_i$ is infeasible. By the finite intersection property, the whole family has a common point, and we let $R^{\mathfrak{M}}(\bar{w})$ be one. As a point of $\mathcal{V}_R$, this value respects the a priori bound of R, and it equals 1 when R is $\mathbf{1}$. The logic requires of the interpretations $F^{\mathfrak{M}}$ and $R^{\mathfrak{M}}$ only that they have the appropriate sorts and, for $R^{\mathfrak{M}}$, that it respect the a priori bound of R (see Section 2). The interpretations chosen above therefore complete the definition of $\mathfrak{M}$.

We claim that every commitment $\varphi \in I$ of $\mathfrak{p}_\infty$ holds in $\mathfrak{M}$. The claim is proved by induction on the construction of φ and, for atomic sentences, on the number of occurrences of function symbols of the original vocabulary $\mathcal{L}$. For a basic sentence, the claim holds by the choice of $R^{\mathfrak{M}}$. A non-basic atomic sentence $R(\bar{t})$ contains a term $F(\bar{w})$ whose arguments

$\bar{w}$ are witnesses (possibly none, when F is a constant symbol). By fairness, Olivia made the challenge (Tm) for the commitment $R(\bar{t}) \in I$ and this term, and Josh added $R(\bar{t}') \in I$, which holds by the induction hypothesis. The sentences $R(\bar{t})$ and $R(\bar{t}')$ have the same value in $\mathfrak{M}$, because the witness w' that replaces $F(\bar{w})$ is $F^{\mathfrak{M}}(\bar{w})$. For $f(\varphi_1, \dots, \varphi_\ell) \in I$, fairness gives, for every $k \in \mathbb{N}$, intervals $D_1, \dots, D_\ell$ such that $\mathfrak{p}_\infty$ contains the commitments $\varphi_j \in D_j \cap I_{\varphi_j}$ and the connective f maps $(D_1 \cap I_{\varphi_1}) \times \cdots \times (D_\ell \cap I_{\varphi_\ell})$ into $I^{(k)}$. By the induction hypothesis, the value of $f(\varphi_1, \dots, \varphi_\ell)$ in $\mathfrak{M}$ lies in $I^{(k)}$ for every k, hence in I. For $\sup_x \varphi(x) \in [\alpha, \beta]$, fairness adds to $\mathfrak{p}_\infty$ a commitment $\varphi(w) \in [-\bar{r}_\varphi, \beta]$ for every witness w of the sort of x, that is, for every element of $\mathfrak{M}$ of this sort. Fairness also adds to $\mathfrak{p}_\infty$, for every $k \in \mathbb{N}$, a commitment $\varphi(w) \in [\alpha - 2^{-k}, \bar{r}_\varphi]$ for some witness w. By the induction hypothesis, the supremum lies in $[\alpha, \beta]$. The remaining cases of the induction are similar. Finally, every axiom $\psi \le r$ of T was named in a challenge (Ax), to which Josh answered with the commitment $\psi \in [\min(-\bar{r}_\psi, r), r]$. By the claim, the structure $\mathfrak{M}$ satisfies $\psi \le r$; the axioms $\psi \ge r$ are treated in the same way.

Hence the reduct of $\mathfrak{M}$ to $\mathcal{L}$ is a model of T.

(2) and (3) are equivalent. Regard $\mathrm{G}(T)$ as an infinite game of perfect information in which a move that the rules do not allow loses for the player who makes it, and in which the play continues with arbitrary moves once a player has lost. Olivia then wins exactly when Josh loses at some finite round before she has broken the rules. We give the space of infinite sequences of moves the topology whose basic open sets consist of the sequences that extend a given finite one; since each of Olivia's wins is decided at a finite round, the set of her wins is *open* in this topology. An infinite game of perfect information is *determined* when one of the two players has a winning strategy, and the Gale–Stewart theorem [GS53] asserts that such a game is determined whenever the set of wins of one of the players is open. Since a winning strategy in this game breaks no rule before the opponent has lost, it restricts to a winning strategy in $\mathrm{G}(T)$; conversely, a winning strategy in $\mathrm{G}(T)$, extended by arbitrary moves, wins this game. Hence (3) implies (2), and since the two players cannot both have winning strategies, (2) implies (3). □

Remark 3.4 (Baire category and the Banach–Mazur game). Olivia's bookkeeping can be read as a Baire category argument in the space of plays, and as a winning strategy in a Banach–Mazur game on the same space. We set out both readings in Section B, beside the Baire category argument of Aubone, Dueñez and Iovino [ADI26] in the space of Henkin types.

3.3. **Finite obstructions.** When a theory has no models, the model existence theorem supplies a winning strategy whose tree of plays organizes the resulting finite contradictions. Let σ be a strategy for Olivia. The *tree of* σ, denoted $\mathcal{P}_\sigma$, is the set of finite plays in which Olivia follows σ and Josh has not lost before the last round, ordered by extension. A finite play of $\mathcal{P}_\sigma$ in which Josh has not lost has as children its extensions by Olivia's next challenge and a legal answer of Josh. When Josh has no legal answer to this challenge, the only child of the play is its extension by the challenge alone. A *leaf* of $\mathcal{P}_\sigma$ is a finite play with no children, that is, one at which Josh has lost. The strategy σ is winning exactly when $\mathcal{P}_\sigma$ has no infinite branch, that is, when $\mathcal{P}_\sigma$ is *well-founded*. The *rank* of a well-founded tree is defined by recursion: a leaf has rank 0, and every other node has as rank the least ordinal exceeding the ranks of all its children. The rank of the tree is the rank of its root. The tree $\mathcal{P}_\sigma$ is countably branching, hence countable; when it is well-founded, its rank is therefore a

countable ordinal. This rank is finite exactly when the plays in which Olivia follows σ have bounded length.

To distinguish the countable choices from the intervening finite choices, we retain only Josh's answers to countable challenges. In Corollary 3.5 below we show that only countable challenges can give a node of $\mathcal{P}_\sigma$ infinitely many children. The *countable record* of a node of $\mathcal{P}_\sigma$ is the finite sequence of the numbers n that Josh names in his countable answers along the branch from the root to the node. The numbers appear in this sequence in the order in which Josh names them. The *skeleton* of σ is the set of countable records of nodes of $\mathcal{P}_\sigma$, ordered by extension; it is a tree of finite sequences of natural numbers.

Corollary 3.5 (Finite obstructions). *Let T be a nonempty countable theory of real-valued $L_{\omega_1\omega}$ whose axioms are closed conditions $\psi \le r$ or $\psi \ge r$ for sentences ψ and rational numbers r, and suppose that T has no models. Then Olivia has a winning strategy in $\mathrm{G}(T)$, and every winning strategy σ of Olivia has the following properties:*

(1) The tree $\mathcal{P}_\sigma$ is well-founded: every play in which Olivia follows σ ends after finitely many rounds, at a leaf.

(2) At every leaf, the last position $\mathfrak{p}$ is a finite set of commitments and records, involving finitely many axioms of T and finitely many witnesses, that is satisfied in no structure under any interpretation of the witnesses. Either $\mathfrak{p}$ is infeasible, in which case a basic sentence $R(\bar{w})$ receives finitely many rational intervals with no common point in $\mathcal{V}_R$, or Josh has no legal answer to Olivia's last challenge, which names a commitment of $\mathfrak{p}$ that holds in no structure.

(3) A node of $\mathcal{P}_\sigma$ has infinitely many children only if σ makes a countable challenge at it; every other node has finitely many children.

(4) If every node of $\mathcal{P}_\sigma$ has finitely many children, then the tree $\mathcal{P}_\sigma$ is finite. In that case there is a number N_σ such that every play in which Olivia follows σ ends within N_σ rounds, and some finite set of axioms of T has no models. Every node of $\mathcal{P}_\sigma$ has finitely many children when no axiom of T contains a countable supremum or infimum.

(5) The skeleton of σ is well-founded, and for every element $\vec{n}$ of the skeleton there is a number $N_\sigma(\vec{n})$ such that every play in which Olivia follows σ and whose countable record is $\vec{n}$ ends within $N_\sigma(\vec{n})$ rounds.

Proof. By Theorem 3.3, Olivia has a winning strategy; let σ be any winning strategy of Olivia. Part (1) holds because σ is winning. For part (2), let $\mathfrak{p}$ be the last position of a leaf, a finite set of commitments and records like every position. If Josh lost because his last answer made $\mathfrak{p}$ infeasible, then the position $\mathfrak{p}$ is satisfied in no structure, since the value of a basic sentence $R(\bar{w})$ in a structure lies in $\mathcal{V}_R$. Otherwise Josh had no legal answer to Olivia's last challenge, which therefore is not a challenge (Ax), (Tot) or (Tm), because under these rules a legal answer always exists. The challenge names a commitment of $\mathfrak{p}$, and the argument for "(1) implies (2)" in the proof of Theorem 3.3, which uses the axioms of T only under (Ax), produces a legal answer whenever this commitment holds in some structure under some interpretation of its witnesses. The commitment therefore holds in no structure, and neither does $\mathfrak{p}$. Part (3) follows from the count of Josh's possible answers in the paragraph after Definition 3.2.

Parts (4) and (5) rest on Kőnig's lemma: an infinite tree in which every node has finitely many children has an infinite branch. For part (4), the tree $\mathcal{P}_\sigma$ is well-founded and by hypothesis finitely branching; hence it is finite, and the length of its longest branch gives N_σ. By part (3), the hypothesis holds when σ makes no countable challenge. In particular, the

hypothesis holds when no axiom of T contains a countable supremum or infimum. Indeed, every sentence that occurs in a play is atomic or is an instance of a subformula of an axiom. It remains to be shown that some finite set of axioms of T has no models. Let T_σ be the set of axioms named in challenges (Ax) at nodes of $\mathcal{P}_\sigma$, together with one axiom of T if there are none. In $\mathrm{G}(T_\sigma)$, every challenge that σ makes at a node of $\mathcal{P}_\sigma$ is allowed by the rules, and Josh has the same answers as in $\mathrm{G}(T)$. The strategy σ therefore wins $\mathrm{G}(T_\sigma)$, and the finite theory T_σ has no models by Theorem 3.3.

For part (5), let $\vec{n}$ be a finite sequence of natural numbers, and let $\mathcal{P}_\sigma^{\vec{n}}$ be the set of nodes of $\mathcal{P}_\sigma$ whose countable record is an initial segment of $\vec{n}$. The set $\mathcal{P}_\sigma^{\vec{n}}$ is a subtree of $\mathcal{P}_\sigma$, and every play in which Olivia follows σ and whose countable record is $\vec{n}$ runs inside this subtree. In $\mathcal{P}_\sigma^{\vec{n}}$, a node at which σ makes a countable challenge has at most one child. If Josh has no legal answer to the challenge, its only child is the extension by the challenge alone. Otherwise, if the countable record of the node is $\vec{n}$ itself, the node has no child in $\mathcal{P}_\sigma^{\vec{n}}$. In the remaining case, the only child of the node is the extension by the countable answer whose number is the next entry of $\vec{n}$. By part (3), every other node of $\mathcal{P}_\sigma^{\vec{n}}$ has finitely many children. Since $\mathcal{P}_\sigma^{\vec{n}}$ is well-founded, it is finite by Kőnig's lemma, and the length of its longest branch gives $N_\sigma(\vec{n})$. Finally, suppose, towards a contradiction, that the skeleton has an infinite branch $\vec{n}_1 \subsetneq \vec{n}_2 \subsetneq \cdots$. The set of nodes of $\mathcal{P}_\sigma$ whose countable record is an initial segment of some $\vec{n}_i$ then forms a finitely branching subtree of $\mathcal{P}_\sigma$. This subtree contains nodes of every countable record $\vec{n}_i$ and is therefore infinite. By Kőnig's lemma, it has an infinite branch, contradicting part (1); the skeleton is therefore well-founded. □

Part (4) of Corollary 3.5 recovers compactness for countable theories of closed conditions without countable suprema and infima: if such a theory has no models, some finite set of its axioms already has no models. For the infinitary logic, part (3) identifies the additional choices that may prevent this conclusion, and part (5) bounds the lengths of plays with a fixed countable record (see Figures 5 and 6). Parts (1) and (2) supply a well-founded tree of finite obstructions, one at each leaf, even when no finite set of axioms is inconsistent. The extension of the theory of Theorem B by the requirement that μ take only the values 0 and 1 gives an example. This is the theory T^{01} of the Introduction; in Proposition 8.3 below we show that every finite set of its axioms has a model, although the whole theory has no models by Theorem C.

Proof of Theorem A. In the model existence game, Olivia is the player who claims that the theory has no models, and the tree of plays of a strategy σ of hers is the tree $\mathcal{P}_\sigma$. Let T' be as in Theorem A, with its axioms written in the normal form of Subsection 3.1. By Theorem 3.3, the theory T' has no models if and only if Olivia has a winning strategy in $\mathrm{G}(T')$. If T' has no models, the remaining assertions of Theorem A are those of Corollary 3.5, with T' in place of T. At a leaf, the commitments of the last position are satisfied in no structure even when its records are disregarded, as Theorem A asserts: by part (2) of the corollary, either they assign to a basic sentence $R(\bar{w})$ intervals with no common point in $\mathcal{V}_R$, or one of them holds in no structure. □

Remark 3.6 (Strict conditions). The conditions of the logic are strict inequalities, and a closed condition is a countable theory of strict ones (see Section 2). For a theory that contains strict conditions, the game has one more rule.

(Str): Olivia names a condition $\varphi < r$ of T, for a sentence φ. Josh chooses a rational number $\varrho < r$ and adds $\varphi \in [-\bar{r}_\varphi, \varrho]$. For a condition $\varphi > r$ of T, Josh chooses a rational number $\varrho > r$ and adds $\varphi \in [\varrho, \bar{r}_\varphi]$.

With one more case for this rule, the proof of Theorem 3.3 applies to theories that consist of closed conditions as above and of strict conditions on sentences. In the additional case, a truthful Josh chooses ϱ between the value of φ and r, and in the structure built from a fair play the commitment $\varphi \in [-\bar{r}_\varphi, \varrho]$ gives $\varphi \leq \varrho < r$. Under (Str), however, Josh chooses among infinitely many answers, and the challenges (Str) are accordingly counted as countable challenges in Corollary 3.5. In the countable record, Josh's rational answers are coded by natural numbers; parts (1), (3) and (5) of the corollary then hold. Part (2) holds once the condition named in the last challenge is added to $\mathfrak{p}$. Indeed, Josh has no legal answer to (Str) for $\varphi < r$ exactly when $r \leq -\bar{r}_\varphi$, and none for $\varphi > r$ exactly when $r \geq \bar{r}_\varphi$. In either case, the condition named is itself unsatisfiable. With the same counting, part (4) holds, except for its clause on theories without countable suprema and infima. The exception arises because a challenge (Str) with a legal answer has infinitely many children and because such theories may have no models although each finite set of their axioms has one. For a nullary predicate symbol R with $r_R > 0$, the theory $\{R > 0\} \cup \{R < \frac{1}{n} : n \geq 1\}$ is such a theory. Olivia wins by naming first $R > 0$, to which Josh answers with some rational $\varrho > 0$ and the commitment $R \in [\varrho, \bar{r}_R]$, and then a condition $R < \frac{1}{n}$ with $\frac{1}{n} < \varrho$, to which every legal answer of Josh adds an interval disjoint from $[\varrho, \bar{r}_R]$ and thereby makes the position infeasible. The tree of this strategy branches infinitely at its root. The contrast between the two kinds of conditions has a topological counterpart in the space of Henkin types of Aubone, Dueñez and Iovino [ADI26], which we describe in Section B.

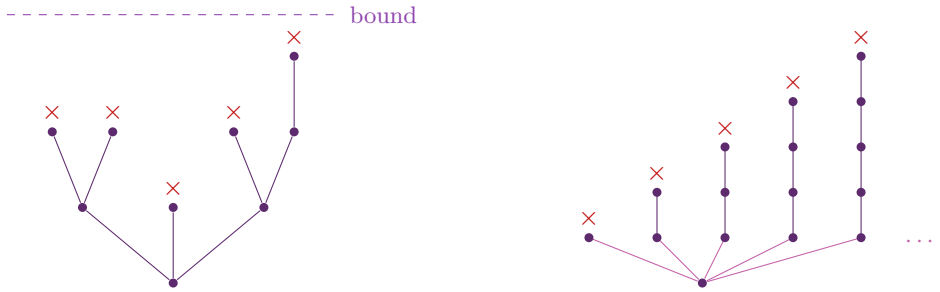


FIGURE 5. Trees growing upward from their roots: a finitely branching tree without infinite branches, which is finite (left), and a countably branching tree without infinite branches, with one branch of each finite length (right).

Historical note. The Banach–Mazur game began as Problem 43 of the Scottish Book, posed by Mazur. On 4 August 1935 Banach wrote in the book that Mazur's conjecture is true. His proof was never published; Oxtoby published a proof in 1957 [Mau81, Oxt57]. Gale and Stewart proved in 1953 that open and closed games are determined [GS53]. The analysis of finite obstructions rests, in addition, on the lemma on finitely branching trees that Dénes Kőnig had published in 1927 [Kőn27]. The witnesses of the game descend from the new constants of Henkin's completeness proof of 1949 [Hen49]. For Väänänen there are essentially three kinds of games in logic, the semantic game, the model existence game and the Ehrenfeucht–Fraïssé game; his book of 2011 is built on them [Vää11].

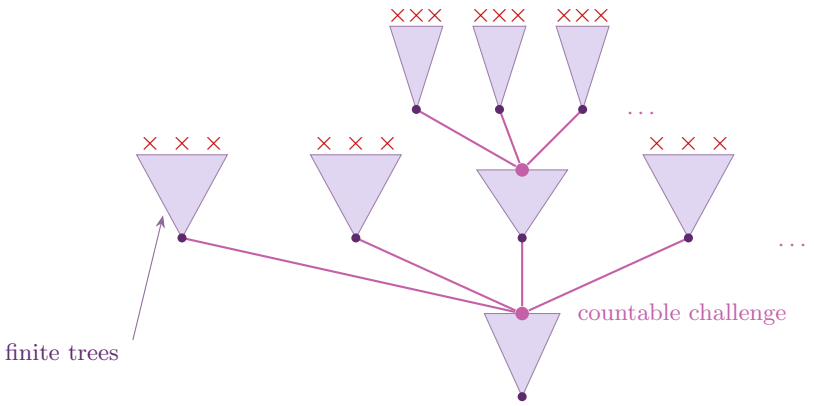


FIGURE 6. The tree of a winning strategy for Olivia, growing upward from its root, with an infinite fan at each countable challenge to which Josh has a legal answer and finite trees between countable challenges.

4. The vocabulary

We now return to the theory T of quantum probability of Theorem B and present its vocabulary, in which the three sorts of Section 1 are linked by projection, coordinate and span maps.

The vocabulary $\mathcal{L}$ has exactly three sorts, the Hilbert sort, the Solèr sort and the sequence sort. The Hilbert and sequence sorts are vector sorts, intended to be the closed unit balls of complex Hilbert spaces; this restriction keeps their norm and inner-product predicates bounded. In the intended interpretation, the elements of the Solèr sort name closed subspaces and the span function assigns a lattice element to a sequence of vectors. The list below specifies the symbols and their individual a priori bounds, while the axioms govern their interaction and standardness remains a separate hypothesis.

- The *Hilbert sort* $\mathbf{H}$ has a binary function $(x, y) \mapsto \lambda x + \lambda' y$ for each pair λ, λ' of elements of the field $\mathbb{Q}(i)$ of complex numbers with rational real and imaginary parts such that $|\lambda| + |\lambda'| \leq 1$. The sort $\mathbf{H}$ also has the predicates $\mathrm{Re}\langle x \mid y\rangle$, $\mathrm{Im}\langle x \mid y\rangle$ and $\|x\|$, with a priori bound 1, and constants e_n for $n \in \mathbb{N}$. We write 0 for the closed term $0e_0 + 0e_0$, and $\langle x \mid y\rangle$ for the pair of predicates $\mathrm{Re}\langle x \mid y\rangle$ and $\mathrm{Im}\langle x \mid y\rangle$. We also write $|\langle x \mid y\rangle|^2$ for the formula $(\mathrm{Re}\langle x \mid y\rangle)^2 + (\mathrm{Im}\langle x \mid y\rangle)^2$, built with the squares of Lemma 2.1.
- The *sequence sort* $\mathbf{K}$ has the same functions and predicates as $\mathbf{H}$, one constant ε, and unary functions res_N for $N \in \mathbb{N}$; in this sort, we write 0 for the closed term $0\varepsilon + 0\varepsilon$.
- The *Solèr sort* $\mathbf{L}$ has the constants 0 and 1, the binary functions $\wedge$ and $\vee$ and the unary function $^\perp$.
- Function symbols connecting the sorts:

$$\pi\colon \mathbf{L} \times \mathbf{H} \to \mathbf{H}, \quad \mathrm{at}\colon \mathbf{H} \to \mathbf{L}, \quad \mathrm{rep}\colon \mathbf{L} \to \mathbf{H},$$
$$\iota_n\colon \mathbf{H} \to \mathbf{K}, \quad \kappa_n\colon \mathbf{K} \to \mathbf{H}\ (n \in \mathbb{N}), \quad \mathrm{span}\colon \mathbf{K} \to \mathbf{L}.$$

- A predicate μ on $\mathbf{L}$, the *measure*, with a priori bound 2. The bound is wider than 1 so that the inequality $\mu(a) \leq 1$, like $\mu(a) \geq 0$, is imposed by the axiom (M1) of Section 5 below, and not by the vocabulary.

- *Truncations.* For every predicate symbol R listed above, every positive integer n and every rational number q, two predicate symbols $R^+_{n,q}$ and $R^-_{n,q}$ with the arguments of R and with a priori bound 1.

The intended values of the truncations are

$$R^+_{n,q}(\bar{x}) = \min\bigl(1, \max(0, n(R(\bar{x}) - q))\bigr), \qquad R^-_{n,q}(\bar{x}) = \min\bigl(1, \max(0, n(q - R(\bar{x})))\bigr),$$

which the axioms (Tr) of Section 5 impose. For an atomic formula $R(\bar{t})$ whose predicate symbol is one of the inner-product, norm or measure predicates listed above, and a rational number q, let

$$\lceil R(\bar{t}) > q \rceil = \bigvee_{n \geq 1} R^+_{n,q}(\bar{t}), \qquad \lceil R(\bar{t}) < q \rceil = \bigvee_{n \geq 1} R^-_{n,q}(\bar{t}),$$

countable joins of members with a priori bound 1. In every structure that satisfies (Tr), the formula $\lceil R(\bar{t}) > q \rceil$ takes the value 1 where $R(\bar{t}) > q$ and the value 0 elsewhere, and likewise for $\lceil R(\bar{t}) < q \rceil$. Several symbols now have more than one meaning: the symbol 0 names a sentence, a closed term of each vector sort and a constant of the Solèr sort, and the symbol 1 names a sentence and a constant of the Solèr sort. The symbols $\wedge$ and $\vee$, for their part, name both connectives and functions of the Solèr sort, and the arguments to which they are applied always show which is meant. For vector terms t_1, t_2 of the same sort, the equation $t_1 \doteq t_2$ abbreviates the axiom $\|\frac{1}{2}t_1 - \frac{1}{2}t_2\| \leq 0$, a closed condition when t_1 and t_2 are closed. A vector x *belongs* to an element a of the Solèr sort, or is a *member* of a, when $\pi(a, x) \doteq x$.

In the intended interpretation, the connecting symbols have the following roles. Here a is a lattice element, x a vector, s a sequence, and $n, N \in \mathbb{N}$.

Symbol	Intended meaning
$\pi(a, x)$	Orthogonal projection of x onto the subspace named by a.
$\mathrm{at}(x)$	The ray spanned by x if $x \neq 0$; the zero element otherwise.
$\mathrm{rep}(a)$	A chosen unit vector in a nonzero subspace a; the zero vector for $a = 0$.
$\iota_n x$	The sequence with x in position n and zeros elsewhere.
$\kappa_n s$	The entry of s in position n.
$\mathrm{res}_N s$	The sequence obtained by keeping the entries with indices $n < N$ and replacing the others by zero.
$\mathrm{span}(s)$	The lattice element naming the closed span of the entries of s.

The constants $(e_n)_{n \in \mathbb{N}}$ name an orthonormal basis of the Hilbert space. The sequence sort consists of the square-summable sequences of vectors of the Hilbert space with ℓ^2 norm at most 1, and ε names $(2^{-n-1}e_n)_{n \in \mathbb{N}}$. The named sequence ε, together with the axioms (K4) and (M5) of Section 5, is used in the proof of Theorem C to locate a basis atom of measure 1 when μ takes only the values 0 and 1. The parity argument of that proof then takes place in the span of a block of four consecutive basis vectors, chosen to contain this atom.

The language has no primitive equality symbol. For vectors, we already use zero norm of a scaled difference. For lattice elements, the representative vector supplies a test for zero: in the intended interpretation, $\|\mathrm{rep}(a)\|$ is 0 for the zero subspace and 1 otherwise. To test whether $a \leq b$, apply this test to $a \wedge (a \wedge b)^\perp$, the orthogonal complement of $a \wedge b$ in a. This element is zero exactly when $a \leq b$. We therefore define

$$[a \neq 0] = 2\|\mathrm{rep}(a)\| - 1, \qquad [a \leq b] = 1 - 2\|\mathrm{rep}(a \wedge (a \wedge b)^\perp)\|, \qquad [a = b] = [a \leq b] \wedge [b \leq a].$$

In the intended interpretation, these are two-valued tests for nonzero subspace, inclusion and equality of subspaces. In Lemma 5.3 below we associate a closed subspace of a Hilbert space with each element of the Solèr sort in every model of T_0. In Lemma 5.4 we show that the three formulas are two-valued in every such model: their value 1 means, respectively, that the associated subspace is nonzero, that the subspace associated with a is contained in the one associated with b, and that the two subspaces coincide. The equality test does not assert literal identity of the elements of the Solèr sort. For terms t_1, t_2 of the Solèr sort, the equation $t_1 = t_2$ abbreviates the axiom $[t_1 = t_2] \geq 1$.

The axiom of Proposition 7.5 excludes the points of a Borel set from the range of the measure. This axiom needs a two-valued formula that says whether a value of the measure lies in the set. For a Borel set $B \subseteq \mathbb{R}$, we write $\beta_B(a)$ for any formula of real-valued $L_{\omega_1\omega}$, with the single free variable a of the Solèr sort, whose values are prescribed as follows. At every element a of every structure for $\mathcal{L}$ that satisfies the axioms (Tr), the formula takes the value 1 if $\mu(a) \in B$ and -1 otherwise.

Lemma 4.1. *For every Borel set $B \subseteq \mathbb{R}$, there is a formula $\beta_B(a)$ as just described.*

Proof. Let $\mathcal{S}$ be the class of sets $B \subseteq \mathbb{R}$ for which such a formula exists with a priori bound at most 3. For rational numbers $q \leq q'$, the formula

$$1 - 2\big(\lceil \mu(a) > q' \rceil \vee \lceil \mu(a) < q \rceil\big)$$

takes the value 1 when $\mu(a) \in [q, q']$ and -1 otherwise, in every structure that satisfies (Tr). Its a priori bound is 3, and the class $\mathcal{S}$ therefore contains the closed intervals with rational endpoints. If $B \in \mathcal{S}$, then the negation $-\beta_B(a)$, which has the same a priori bound, is a formula for $\mathbb{R} \setminus B$. If $B_n \in \mathcal{S}$ for every $n \in \mathbb{N}$, then the countable join of the formulas $\beta_{B_n}(a)$, whose a priori bounds are at most 3, is a formula for $\bigcup_n B_n$ with a priori bound at most 3. The class $\mathcal{S}$ is therefore a σ-algebra, and since every open subset of $\mathbb{R}$ is a countable union of closed intervals with rational endpoints, it contains every Borel set. □

When a formula $\beta_B(a)$ is obtained by following, as in the proof of Lemma 4.1, a construction of B from closed intervals with rational endpoints by complements and countable unions, its negations and countable joins are nested as in the construction. The depth of this nesting is a countable ordinal α, and an induction on α shows that the set B then belongs to the class $\mathbf{\Delta}^0_{\alpha+2}$ of the Borel hierarchy. Since this hierarchy is proper on $\mathbb{R}$ [Kec95, Section 22], the least depth of such a construction of B is unbounded below ω_1 as B varies. The axioms of Proposition 7.5, with formulas $\beta_B(a)$ built in this way, therefore involve formulas whose joins are nested to arbitrarily large countable depth, although the formulas $\beta_B(a)$ contain no quantifier.

Definition 4.2. Let $\mathfrak{M}$ be a structure for $\mathcal{L}$, or for a vocabulary obtained from $\mathcal{L}$ by removing the predicate μ, the sequence sort with the symbols that involve it, or both, the truncations of the removed predicates being removed with them. The structure $\mathfrak{M}$ is *standard* when its Hilbert sort is the closed unit ball of a complex Hilbert space H, when its sequence sort, if it has one, is the closed unit ball of a complex Hilbert space K, and when the functions $\lambda x + \lambda' y$ and the predicates $\mathrm{Re}\langle x \mid y\rangle$, $\mathrm{Im}\langle x \mid y\rangle$ and $\|x\|$ are the restrictions of the operations, the inner products and the norms of these spaces.

The theories introduced below satisfy $T_0 \subseteq T_1 \subseteq T$: the theory T_0 uses the Hilbert and Solèr sorts, T_1 adds sequences and their spans, and T adds the measure. Standardness concerns the two vector sorts and their Hilbert-space operations. It does not by itself identify the Solèr sort with the full lattice of closed subspaces, or the sequence sort with all square-summable sequences of vectors of the Hilbert sort; these identifications are the results of Sections 6

and 7. Standardness is a hypothesis on the model, not a consequence of the axioms: every standard model of T_0 has an uncountable Hilbert sort (see Proposition 6.1 below), whereas the downward Löwenheim–Skolem theorem gives T_0, T_1 and T models whose sorts are countable (see Remark 10.3). Accordingly, Theorem C concerns every model of T, whereas part (1) of Theorem B and Proposition 7.5 concern the standard models.

5. The axioms

The axioms build the intended structure in three stages: T_0 links the Hilbert and Solèr sorts by projection symbols, T_1 adds represented spans of sequences, and T adds the measure. In standard models, these stages yield a pre-Solèr lattice, the full lattice of closed subspaces, and a measure on the full lattice, in each case taking the quotient by the formula expressing equality.

The axiom groups follow these tasks. Groups (H) and (OL) specify the vector and lattice operations; (P) and (Sub) identify the projections and their lattice operations; (At) links vectors with rays and chooses representatives of nonzero subspaces. Groups (K) and (Sp) supply sequences and their spans, while (M) governs the measure. Group (Tr) fixes the truncations used for exact threshold tests. The formal definitions of T_0, T_1 and T follow the list.

The variables x, y, z range over $\mathbf{H}$, the variables s, t over $\mathbf{K}$, and the variables a, b, c over $\mathbf{L}$; every axiom is a closed condition required of all values of its free variables, as in Section 2. By contrast, the letters λ, λ' and n, m, k, N are metavariables, not variables of the logic: the pair λ, λ' ranges over $\mathbb{Q}(i)$ with $|\lambda| + |\lambda'| \leq 1$, and the indices n, m, k, N range over $\mathbb{N}$. A free occurrence of one of these metavariables gives an axiom schema, one axiom for each permitted choice, whereas bound occurrences index the indicated sums, joins or infima.

Vector equations written with $\doteq$ and lattice equations written with $=$ retain the meanings defined in Section 4. For two-valued formulas θ, θ', the axiom $-\theta \vee \theta' \geq 1$ expresses $\theta \Rightarrow \theta'$. For a two-valued formula θ and any formula ψ, the axiom $\theta \wedge \psi \leq 0$ expresses that θ implies $\psi \leq 0$. Italic readings describe the geometry of the standard models. For $N \in \mathbb{N}$, write $b_N = \mathrm{at}(e_0) \vee \cdots \vee \mathrm{at}(e_{N-1})$, with $b_0 = 0$.

(H) The Hilbert sort.

(H1) *Inner product space.* For any two terms built from variables by the functions $\lambda x + \lambda' y$ that denote the same $\mathbb{Q}(i)$-linear combination of the variables, the equation between them; the equations $\langle \lambda x + \lambda' y \mid z \rangle = \lambda \langle x \mid z \rangle + \lambda' \langle y \mid z \rangle$ and $\langle y \mid x \rangle = \overline{\langle x \mid y \rangle}$, each written as two real equations; and $\|x\| \geq 0$ and $\|x\|^2 = \mathrm{Re}\langle x \mid x \rangle$.

(H2) *Separability.* $\langle e_n \mid e_m \rangle = \delta_{nm}$, where δ_{nm} is the Kronecker delta, and

$$\bigwedge_N \left\| \tfrac{1}{2} x - \tfrac{1}{2} \pi(b_N, x) \right\| \leq 0.$$

The constants e_n are orthonormal, and every vector is the limit of its projections onto the spans of $e_0, \ldots, e_{N-1}$.

(OL) The Solèr sort.

(OL) The equations of ortholattices:

$$a \wedge b = b \wedge a, \quad a \wedge (b \wedge c) = (a \wedge b) \wedge c, \quad a \wedge (a \vee b) = a,$$
$$a \vee b = b \vee a, \quad a \vee (b \vee c) = (a \vee b) \vee c, \quad a \vee (a \wedge b) = a,$$
$$a \wedge 0 = 0, \quad a \vee 1 = 1, \quad a^{\perp\perp} = a, \quad (a \wedge b)^{\perp} = a^{\perp} \vee b^{\perp}, \quad a \wedge a^{\perp} = 0.$$

(P) Each element of the Solèr sort determines a projection.

(P1) $\pi(a, \lambda x + \lambda' y) \doteq \lambda \pi(a, x) + \lambda' \pi(a, y)$.
(P2) $\pi(a, \pi(a, x)) \doteq \pi(a, x)$.
(P3) $\langle \pi(a, x) \mid y \rangle = \langle x \mid \pi(a, y) \rangle$.

The projection axioms (P) concern each lattice element separately, while group (Sub) relates the projection maps to the lattice operations. For a closed subspace W of a Hilbert space, let P_W denote the orthogonal projection onto W. No noncommutative polynomial with fixed coefficients yields $P_{U \cap V}$ from P_U and P_V for every pair of closed subspaces U and V, since such polynomials are continuous in the norm topology and the map $(P_U, P_V) \mapsto P_{U \cap V}$ is not. The projection $P_{U \cap V}$ is the strong limit of the powers $(P_U P_V)^n$.

Theorem 5.1 (von Neumann's alternating projection theorem [vN49, Lemma 22]**; see also** [vN50]**).** *Let U and V be closed subspaces of a Hilbert space, and let P_U, P_V and $P_{U \cap V}$ be the orthogonal projections onto U, V and $U \cap V$. Then, for every vector x of the space, the sequence $\big((P_U P_V)^n x\big)_{n \in \mathbb{N}}$ converges in norm to $P_{U \cap V} x$.*

Axiom (Sub1) below requires the projection of $a \wedge b$ to be the limit in Theorem 5.1. Axiom (Sub2) requires the projection of $a \vee b$ to be the identity minus the same limit for $a^{\perp}$ and $b^{\perp}$, since the closure of $U + V$ is $(U^{\perp} \cap V^{\perp})^{\perp}$. In Lemma 5.4 the theorem then identifies the subspace represented by $a \wedge b$ with the intersection of the subspaces represented by a and b, and the subspace represented by $a \vee b$ with the closure of their sum. The convergence admits no rate that is uniform in U and V. For $n \geq 1$ the norm of $(P_U P_V)^n - P_{U \cap V}$ is c^{2n-1}, where c is the cosine of the Friedrichs angle between U and V [KW88]. When $U + V$ is not closed the convergence is arbitrarily slow [BDH09]. Accordingly, axioms (Sub1) and (Sub2) assert convergence through the limit superior $\bigwedge_N \bigvee_{n \geq N}$, which vanishes exactly when the nonnegative errors tend to zero.

(Sub) The operations of the Solèr sort are those of subspaces.

(Sub0) $\pi(0, x) \doteq 0$ and $\pi(1, x) \doteq x$.
Only the zero vector belongs to 0; every vector belongs to 1.

(Sub1) $\bigwedge_N \bigvee_{n \geq N} \left\| \frac{1}{2} \pi(a \wedge b, x) - \frac{1}{2} (P_a P_b)^n x \right\| \leq 0$, where $(P_a P_b)^n x$ is the term obtained by applying $\pi(b, \cdot)$ and then $\pi(a, \cdot)$, n times.
The members of $a \wedge b$ are the common members of a and b.

(Sub2) $\bigwedge_N \bigvee_{n \geq N} \left\| \frac{1}{4} \pi(a \vee b, x) - \frac{1}{4} x + \frac{1}{4} (P_{a^{\perp}} P_{b^{\perp}})^n x \right\| \leq 0$, where $(P_{a^{\perp}} P_{b^{\perp}})^n x$ is the term of (Sub1) for $a^{\perp}$ and $b^{\perp}$, and the vector inside the norm is denoted by a term built from the functions $\lambda x + \lambda' y$.
The closed subspace spanned by the members of $a \vee b$ is the closure of the sum of the closed subspaces spanned by the members of a and by those of b.

(Sub3) $\frac{1}{2}\pi(a^\perp, x) \doteq \frac{1}{2}x - \frac{1}{2}\pi(a, x)$.
The members of $a^\perp$ are the vectors orthogonal to all members of a.

(At) Atoms and unit vectors.

(At1) $\pi(\mathrm{at}(x), x) \doteq x$, and $\|\pi(\mathrm{at}(x), y)\|^2 \, \|x\|^2 = |\langle y \mid x\rangle|^2$.
Every vector x belongs to $\mathrm{at}(x)$, and for $x \neq 0$ the projection onto $\mathrm{at}(x)$ is the projection onto the line spanned by x.

(At2) $\pi(a, \mathrm{rep}(a)) \doteq \mathrm{rep}(a)$; $\|\mathrm{rep}(a)\| \wedge \big(1 - \|\mathrm{rep}(a)\|\big) \leq 0$; and $\|\pi(a, x)\| \leq \|\mathrm{rep}(a)\|$.
The vector $\mathrm{rep}(a)$ belongs to a; its norm is 0 or 1; and it is 1 unless a has no nonzero member.

(At3) $\|\mathrm{rep}(\mathrm{at}(x))\| \leq \lceil \|x\| > 0 \rceil$.
If $\|x\| = 0$, then $\mathrm{at}(x)$ has no nonzero member.

(K) The sequence sort.

(K1) The axioms (H1), written for **K**.

(K2) $\iota_n(\lambda x + \lambda' y) \doteq \lambda \iota_n x + \lambda' \iota_n y$ and $\kappa_n(\lambda s + \lambda' t) \doteq \lambda \kappa_n s + \lambda' \kappa_n t$; $\langle \iota_n x \mid \iota_m y\rangle = \delta_{nm}\langle x \mid y\rangle$; and $\langle \kappa_n s \mid x\rangle = \langle s \mid \iota_n x\rangle$.

(K3) $\langle \mathrm{res}_N s \mid t\rangle = \sum_{n<N} \langle \kappa_n s \mid \kappa_n t\rangle$, written as two real equations, and $\bigwedge_N \big\| \frac{1}{2}s - \frac{1}{2}\mathrm{res}_N s \big\| \leq 0$.
The sequence $\mathrm{res}_N s$ keeps the first N entries of s, and these sequences converge to s as N grows.

(K4) $\kappa_n \varepsilon \doteq 2^{-n-1} e_n$.

(Sp) Spans.

(Sp1) $\pi(\mathrm{span}(s), \kappa_n s) \doteq \kappa_n s$.
Every entry of s belongs to $\mathrm{span}(s)$.

(Sp2) $\Big(-\bigwedge_n [\mathrm{at}(\kappa_n s) \le b]\Big) \vee [\mathrm{span}(s) \le b] \ge 1$.
If every entry of s belongs to b, then the span $\mathrm{span}(s)$ *lies below b.*

(Tr) Truncations.

(Tr) For every predicate symbol R of $\mathcal{L}$ other than $\mathbf{1}$ and the truncation symbols, every positive integer n and every rational number q:

$$R^+_{n,q}(\bar{x}) = \big((n \cdot (R(\bar{x}) - q)) \wedge 1\big) \vee 0, \qquad R^-_{n,q}(\bar{x}) = \big((n \cdot (q - R(\bar{x}))) \wedge 1\big) \vee 0.$$

We require the predicate μ to take values in $[0, 1]$, to have value 1 at the greatest element of the Solèr sort, and to agree on elements related by the equality formula $[a = b] = 1$. Axioms (M1)–(M3) impose these requirements, while (M4) imposes orthogonal additivity and (M5) makes the measure continuous along the finite joins that exhaust the represented span of a sequence. In (M4), orthogonality is built into $b \wedge a^\perp$; for any c orthogonal to a, the equality $[c = c \wedge a^\perp] = 1$, together with (M3) and compatibility of equality with the lattice operations (see Corollary 5.5 below), gives the usual additive identity. A term built from the lattice operations involves only finitely many elements of the Solèr sort, and the function span supplies the joins of infinite sequences of atoms. In the standard models every element of the Solèr sort is equal, in the sense of $[a = b] = 1$, to the span $\mathrm{span}(s)$ of a sequence s, and axioms (M5) and (M4) then yield σ-additivity with the help of (M1) and (M3) (see the proof of Theorem B). For $N \in \mathbb{N}$, let $c_N(s)$ be the term $\mathrm{at}(\kappa_0 s) \vee \cdots \vee \mathrm{at}(\kappa_{N-1} s)$, with $c_0(s) = 0$.

(M) The measure.

(M1) $\mu(a) \ge 0$ and $\mu(a) \le 1$.

(M2) $\mu(1) = 1$.

(M3) $[a = b] \wedge |\mu(a) - \mu(b)| \le 0$.
Equal elements have equal measures.

(M4) $\mu\big(a \vee (b \wedge a^\perp)\big) = \mu(a) + \mu(b \wedge a^\perp)$.
The measure is additive on orthogonal elements.

(M5) $\mu(\mathrm{span}(s)) = \bigvee_N \mu(c_N(s))$.
The measure of the span of a sequence is the supremum, over N, of the measures of the joins of the atoms of its first N entries.

Definition 5.2. The theory T_0 consists of the axioms (H), (OL), (P), (Sub), (At) and the axioms (Tr) for the predicates of the Hilbert sort, in the vocabulary obtained from $\mathcal{L}$ by removing the sequence sort with its constant ε, the functions ι_n, κ_n, res_N and span, and the predicate μ, together with the truncations of the removed predicates. The theory T_1 consists of T_0 together with (K), (Sp) and the axioms (Tr) for the predicates of the sequence sort, in the vocabulary $\mathcal{L}$ without μ and its truncations; the theory T consists of T_1 together with (M) and the axioms (Tr) for μ.

We now recover the geometry described by these axioms in every model of T_0, without assuming standardness. First we construct a Hilbert space in which the Hilbert-sort operations are represented and the projection symbols act as orthogonal projections. Von Neumann's alternating projection theorem (Theorem 5.1) then shows that the projection of $a \wedge b$ has as range the intersection of the ranges of the projections of a and b, and that the projection of $a \vee b$ has as range the closure of their sum; equality of the represented subspaces will make the quotient of the Solèr sort an ortholattice.

Lemma 5.3. *Let $\mathfrak{M}$ be a model of T_0. There are a complex Hilbert space $\bar{E}$, a subset Y of the closed unit ball of $\bar{E}$ whose $\mathbb{Q}(i)$-linear span E is dense in $\bar{E}$, and a map $x \mapsto \hat{x}$ from the Hilbert sort of $\mathfrak{M}$ onto Y with the following properties. The map sends $\lambda x + \lambda' y$ to $\lambda\hat{x} + \lambda'\hat{y}$, and the predicates $\langle x \mid y\rangle$ and $\|x\|$ to $\langle \hat{x} \mid \hat{y}\rangle$ and $\|\hat{x}\|$. For every element a of the Solèr sort, the map $\hat{x} \mapsto \widehat{\pi(a,x)}$ is well defined and extends uniquely to an orthogonal projection $\bar{P}_a$ of $\bar{E}$. If $\mathfrak{M}$ is standard, then the space $\bar{E}$ is the Hilbert space whose closed unit ball is the Hilbert sort, and the map $x \mapsto \hat{x}$ is the identity.*

Proof. The idea is to recover a vector space from the Hilbert sort by formal multiples. By (H1), the form given by the pair of predicates $\langle x \mid y\rangle$ is hermitian and satisfies $\mathrm{Re}\langle x \mid x\rangle = \|x\|^2 \geq 0$; the Cauchy–Schwarz inequality therefore holds for it. Each pair (k, x), with $k \geq 1$ an integer and x an element of the Hilbert sort, stands for kx. We define the sum of two pairs by $(k, x) + (l, y) = (2kl, \frac{1}{2l}x + \frac{1}{2k}y)$ and, for an element λ of $\mathbb{Q}(i)$, the multiple $\lambda(k, x) = (km, \frac{\lambda}{m}x)$, where m is a positive integer with $m \geq |\lambda|$; the form is given by $\langle (k, x) \mid (l, y)\rangle = kl\langle x \mid y\rangle$. We identify two pairs when the form vanishes on their difference. Axiom (H1) makes the quotient a $\mathbb{Q}(i)$-vector space E with a positive definite hermitian form, and we let $\hat{x}$ be the class of $(1, x)$; since the class of (k, x) is $k\hat{x}$, the space E is the $\mathbb{Q}(i)$-linear span of the set Y of the vectors $\hat{x}$. The completion $\bar{E}$ of E is a complex Hilbert space, because multiplication by the elements of $\mathbb{Q}(i)$, which are dense in $\mathbb{C}$, extends by continuity to multiplication by complex numbers. For an element a of the Solèr sort, consider the assignment Q that sends a pair (k, x) to $(k, \pi(a, x))$; by (P1)–(P3), it respects sums and multiples, it is idempotent and it is symmetric with respect to the form, in each case up to pairs on which the form vanishes. For every pair v, symmetry and idempotence therefore give $\langle Qv \mid Qv\rangle = \langle v \mid Qv\rangle$, which is at most $\langle v \mid v\rangle^{1/2}\langle Qv \mid Qv\rangle^{1/2}$ by the Cauchy–Schwarz inequality, whence $\langle Qv \mid Qv\rangle \leq \langle v \mid v\rangle$. In particular, the assignment Q sends pairs on which the form vanishes to such pairs, and it therefore induces on E a $\mathbb{Q}(i)$-linear, idempotent, symmetric and contractive map Q_a, which extends by continuity to a self-adjoint idempotent $\bar{P}_a$ of $\bar{E}$. In a standard model, the Hilbert sort is the closed unit ball of a Hilbert space H, and the form is positive definite on it. In that case, the map sending the class of (k, x) to kx identifies E with H, since every vector of H is a positive integer multiple of a vector of the ball. As H is complete, it follows that $\bar{E} = H$ and that $x \mapsto \hat{x}$ is the identity of the ball. □

For each element a of the Solèr sort, write S_a for the range of the projection $\bar{P}_a$ just constructed; in a standard model, S_a is the closed subspace spanned by the members of a, as in the Introduction. We must show that the lattice operations act on these ranges as intended and that the formulas of Section 4 detect their nonzeroness, inclusion and equality. The next lemma proves this in every model of T_0, allowing distinct elements of the sort to name the same subspace.

Lemma 5.4. *Let $\mathfrak{M}$ be a model of T_0, and let $\bar{E}$ and $\bar{P}_a$ be as in Lemma 5.3.*

(1) We have $\bar{P}_0 = 0$, $\bar{P}_1 = \mathrm{Id}$ and $\bar{P}_{a^\perp} = \mathrm{Id} - \bar{P}_a$; moreover, $S_{a\wedge b} = S_a \cap S_b$ and $S_{a\vee b} = \overline{S_a + S_b}$.

(2) The norm $\|\mathrm{rep}(a)\|$ *is* 1 *if* $S_a \neq \{0\}$ *and* 0 *otherwise; hence the formula* $[a \neq 0]$ *is two-valued, with the value* 1 *exactly when* $S_a \neq \{0\}$.
(3) The formula $[a \leq b]$ *is two-valued, with the value* 1 *exactly when* $S_a \subseteq S_b$*, and the formula* $[a = b]$ *is two-valued, with the value* 1 *exactly when* $S_a = S_b$.
(4) For every x *with* $\|x\| > 0$*, the subspace* $S_{\mathrm{at}(x)}$ *is* $\mathbb{C}\hat{x}$*, and for every* x *with* $\|x\| = 0$ *it is* $\{0\}$.

Proof. (1) The first three equalities follow from (Sub0) and (Sub3) by continuity, since E is dense in $\bar{E}$. By (Sub1), the projection $\bar{P}_{a\wedge b}$ agrees on the vectors $\hat{x}$, and hence on $\bar{E}$, with the strong limit of the operators $(\bar{P}_a\bar{P}_b)^n$, which is the projection onto $S_a \cap S_b$ by Theorem 5.1. Likewise, axiom (Sub2) and the equality $\bar{P}_{c^\perp} = \mathrm{Id} - \bar{P}_c$ show that $\bar{P}_{a\vee b} = \mathrm{Id} - P_{S_a^\perp \cap S_b^\perp}$, whose range is $(S_a^\perp \cap S_b^\perp)^\perp = \overline{S_a + S_b}$.

(2) By (At2), together with the axiom $\|x\| \geq 0$ of (H1), the norm $\|\mathrm{rep}(a)\|$ is 0 or 1; by (At2) again, the vector $\widehat{\mathrm{rep}(a)}$ lies in S_a, and every x satisfies $\|\bar{P}_a\hat{x}\| \leq \|\mathrm{rep}(a)\|$. If $S_a \neq \{0\}$, then the projection $\bar{P}_a$ does not vanish on the dense subspace E, and hence not on some $\hat{x}$, whence $\|\mathrm{rep}(a)\| > 0$ and so $\|\mathrm{rep}(a)\| = 1$. If $S_a = \{0\}$, then the membership $\widehat{\mathrm{rep}(a)} \in S_a$ gives $\|\mathrm{rep}(a)\| = 0$.

(3) By (1), the subspace $S_{a\wedge(a\wedge b)^\perp}$ is $S_a \cap (S_a \cap S_b)^\perp$, the orthogonal complement of $S_a \cap S_b$ in S_a. It is $\{0\}$ exactly when $S_a \cap S_b = S_a$, that is, when $S_a \subseteq S_b$. By (2), the formula $[a \leq b]$ is therefore two-valued, with the value 1 exactly when $S_a \subseteq S_b$, and the assertion about $[a = b]$ follows.

(4) Let $P = \bar{P}_{\mathrm{at}(x)}$ and suppose that $\|x\| > 0$, whence $\hat{x} \neq 0$. By (At1), the vector $\hat{x}$ lies in $S_{\mathrm{at}(x)}$, and every y satisfies $\|P\hat{y}\|^2\|\hat{x}\|^2 = |\langle \hat{y} \mid \hat{x}\rangle|^2$. Let P_x be the projection onto $\mathbb{C}\hat{x}$, which satisfies $\|P_x\hat{y}\|^2\|\hat{x}\|^2 = |\langle \hat{y} \mid \hat{x}\rangle|^2$ as well; since the vector $\hat{x}$ lies in the range of P, we have $P_x \leq P$. The identity $\|P\hat{y}\|^2 = \|P_x\hat{y}\|^2 + \|(P - P_x)\hat{y}\|^2$ then gives $(P - P_x)\hat{y} = 0$ for every y, whence $P = P_x$ by density. Finally, if $\|x\| = 0$, then (At3) and the axioms (Tr) for the norm give $\|\mathrm{rep}(\mathrm{at}(x))\| = 0$, whence $S_{\mathrm{at}(x)} = \{0\}$ by (2). □

By part (3) of the lemma, the formula $[a = b]$ detects exactly when two elements of the Solèr sort name the same subspace, and the quotient of the sort by this relation therefore inherits the geometry of the subspaces.

Corollary 5.5. *Let* $\mathfrak{M}$ *be a model of* T_0*, and let* $\bar{E}$ *and* S_a *be as in Lemma 5.4. The relation* $[a = b] = 1$ *is an equivalence relation on the Solèr sort that is respected by* $\wedge$*,* $\vee$ *and* $^\perp$*. The quotient* $\mathbf{L}^{\mathfrak{M}}/{=}$ *by this relation is an ortholattice, and the map that sends the class of* a *to* S_a *is an embedding of ortholattices into the lattice* $L(\bar{E})$ *of closed subspaces of* $\bar{E}$ *that carries the order of the quotient to inclusion.*

Proof. Part (3) of Lemma 5.4 shows that the relation $[a = b] = 1$ holds exactly when $S_a = S_b$, and the relation is therefore an equivalence relation. Since, by part (1) of the lemma, the subspaces $S_{a\wedge c}$, $S_{a\vee c}$ and $S_{a^\perp}$ are determined by S_a and S_c, the relation is respected by the operations. The same part (1) shows that the map sending the class of a to S_a preserves the operations; this map is injective by part (3), and its image is therefore a sub-ortholattice of $L(\bar{E})$ isomorphic to the quotient. Finally, the order of the quotient, in which $a \leq b$ means $a \wedge b = a$, corresponds to inclusion, again by part (3). □

In Lemma 5.4 and Corollary 5.5, the axioms (OL) are not used; they hold in every structure that satisfies the other axioms of T_0, by the embedding into $L(\bar{E})$. We keep them, so that the ortholattice structure of the Solèr sort has axioms of its own (see Remark 10.7).

Historical note. In the continuous logic set out by Ben Yaacov, Berenstein, Henson and Usvyatsov in 2008, the unit ball of a Banach space carries the maps $(x, y) \mapsto \alpha x + \beta y$ with $|\alpha| + |\beta| \leq 1$ as functions. A Hilbert space also carries its inner product as a predicate [BYBHU08]. We use the affine operations and inner-product predicates in the real-valued vocabulary of Section 4; their interpretations are governed by the axioms above, without metric or continuity requirements in the logic. Von Neumann's alternating projection theorem, on which axioms (Sub1) and (Sub2) rest, comes from lecture notes that he distributed in mimeographed form in 1933. The notes were printed in 1950 as the second volume of his *Functional Operators* [vN50].

6. Standard models of T_0: pre-Solèr lattices

In a standard model of T_0, whose vocabulary has neither the sequence sort nor the measure, the Hilbert sort is the closed unit ball of a complex Hilbert space H. Since the vocabulary has no equality symbol, distinct elements of the Solèr sort may represent the same closed subspace of H. By Corollary 5.5, the quotient of this sort by the relation $[a = b] = 1$ of Section 4 is an ortholattice, and sending the class of each element a to the closed subspace spanned by the members of a embeds the quotient into the lattice $L(H)$ of closed subspaces of H. We show that the quotient is a pre-Solèr lattice: it satisfies the conditions for a Solèr lattice except, possibly, completeness. The image of the embedding is join-dense and meet-dense in the sense of Section A.1, and the lattice $L(H)$ is therefore isomorphic to the MacNeille completion of the quotient.

The quotient need not be complete, as the standard model of Example 6.5 shows. Nevertheless, its σ-additive states, defined below, are given by density operators (see Theorem 6.6). In Section 7, the sequence sort supplies the remaining subspaces in standard models of T_1.

Fix a standard model $\mathfrak{M}$ of T_0, and let H be the Hilbert space whose closed unit ball is its Hilbert sort (see Definition 4.2). In this case the Hilbert space $\bar{E}$ constructed in Lemma 5.3 is H itself. Write $\mathbf{L}^{\mathfrak{M}}$ for the Solèr sort. For each $a \in \mathbf{L}^{\mathfrak{M}}$, let P_a be the projection $\bar{P}_a$ extending $\pi(a, \cdot)$, and let S_a be its range. Thus S_a is the closed subspace spanned by the members of a, the vectors x of the unit ball with $\pi(a, x) = x$. Write $\tilde{a}$ for the class of a in $\mathbf{L}^{\mathfrak{M}}/{=}$; by Lemma 5.4, two such classes coincide exactly when their subspaces coincide.

Proposition 6.1. *The space H is a separable infinite-dimensional complex Hilbert space, and the constants e_n of the Hilbert sort, for $n \in \mathbb{N}$, form an orthonormal basis of H.*

Proof. By (H2), the vectors e_n are orthonormal. By Lemma 5.4, the projection P_{b_N} is the projection onto the span of $e_0, \ldots, e_{N-1}$, and (H2) places every vector of the unit ball, and hence every vector of H, in the closed span of the vectors e_n. Thus the sequence $(e_n)_{n\in\mathbb{N}}$ is an orthonormal basis of H, which is therefore separable and infinite-dimensional; it is complex by Definition 4.2. □

Proposition 6.2. *The map $\tau\colon \tilde{a} \mapsto S_a$ is an embedding of ortholattices from $\mathbf{L}^{\mathfrak{M}}/{=}$ into the lattice $L(H)$ of closed subspaces of H: it is injective and satisfies*

$$S_0 = \{0\}, \quad S_1 = H, \quad S_{a\wedge b} = S_a \cap S_b, \quad S_{a\vee b} = \overline{S_a + S_b}, \quad S_{a^\perp} = S_a^\perp,$$

and each of the relations $S_a \subseteq S_b$ and $[a \leq b] = 1$ holds exactly when $\tilde{a} \leq \tilde{b}$.

Proof. The proposition restates Corollary 5.5 for the standard model $\mathfrak{M}$, in which $\bar{E} = H$; part (1) of Lemma 5.4 gives the displayed equalities, and part (3) the description of the order. □

Proposition 6.3. *For every nonzero x in the unit ball, the subspace $S_{\mathrm{at}(x)}$ is $\mathbb{C}x$. The atoms of $\mathbf{L}^{\mathfrak{M}}/{=}$ are exactly the classes $\widetilde{\mathrm{at}(x)}$ with $x \neq 0$, and the map $x \mapsto \widetilde{\mathrm{at}(x)}$ induces a bijection from the unit sphere of H, modulo multiplication by complex numbers of modulus one, onto the set of atoms of $\mathbf{L}^{\mathfrak{M}}/{=}$, whose inverse sends an atom $\tilde{a}$ to the class of* $\mathrm{rep}(a)$.

Proof. The first assertion is part (4) of Lemma 5.4, since in a standard model every nonzero vector x has positive norm. Since the ray $\mathbb{C}x$ is an atom of $L(H)$ and the map τ is an order embedding, the class $\widetilde{\mathrm{at}(x)}$ is an atom of $\mathbf{L}^{\mathfrak{M}}/{=}$. Let $\tilde{a}$ be an atom of $\mathbf{L}^{\mathfrak{M}}/{=}$. By part (2) of Lemma 5.4, the vector $\mathrm{rep}(a)$ is a unit vector, and it belongs to a by (At2). Hence the ray $S_{\mathrm{at}(\mathrm{rep}(a))} = \mathbb{C}\mathrm{rep}(a)$ lies in S_a, and the class $\widetilde{\mathrm{at}(\mathrm{rep}(a))}$ is a nonzero element below $\tilde{a}$; since $\tilde{a}$ is an atom, we have $\tilde{a} = \widetilde{\mathrm{at}(\mathrm{rep}(a))}$. For unit vectors x and y, the classes $\widetilde{\mathrm{at}(x)}$ and $\widetilde{\mathrm{at}(y)}$ coincide exactly when $\mathbb{C}x = \mathbb{C}y$, that is, when y is a multiple of x by a complex number of modulus one. In particular, for a unit vector x, the unit vector $\mathrm{rep}(\mathrm{at}(x))$ is such a multiple of x, because it lies in $S_{\mathrm{at}(x)} = \mathbb{C}x$. The inverse of the bijection is therefore as stated. □

Theorem 6.4. *Let $\mathfrak{M}$ be a standard model of T_0, let H be the Hilbert space whose closed unit ball is its Hilbert sort, and, for each element a of its Solèr sort $\mathbf{L}^{\mathfrak{M}}$, let S_a be the closed subspace of H spanned by the members of a. Then the quotient $\mathbf{L}^{\mathfrak{M}}/{=}$ is a pre-Solèr lattice, and the image of the map $\tau\colon \tilde{a} \mapsto S_a$ is join-dense and meet-dense in the lattice $L(H)$ of closed subspaces of H, which is therefore isomorphic to the MacNeille completion of $\mathbf{L}^{\mathfrak{M}}/{=}$. The image of τ contains every finite-dimensional subspace of H and the orthogonal complement of each, and the map τ sends every join that exists in $\mathbf{L}^{\mathfrak{M}}/{=}$ to the corresponding join in $L(H)$.*

Proof. The atoms already represent all rays, so their finite joins give the finite-dimensional subspaces. We use these subspaces to verify that $\mathbf{L}^{\mathfrak{M}}/{=}$ is a pre-Solèr lattice. The zero subspace is S_0; now let F be a nonzero finite-dimensional subspace of H with a basis $x_1, \ldots, x_k$ of unit vectors. By Propositions 6.2 and 6.3, the map τ sends $(\mathrm{at}(x_1) \vee \cdots \vee \mathrm{at}(x_k))^{\sim}$ to the closure of F, which is F itself, and $((\mathrm{at}(x_1) \vee \cdots \vee \mathrm{at}(x_k))^{\perp})^{\sim}$ to $F^{\perp}$.

Condition (OM) holds because $\tau(\mathbf{L}^{\mathfrak{M}}/{=})$ is a sub-ortholattice of the orthomodular lattice $L(H)$. For (A), let $\tilde{a}$ be a nonzero element of $\mathbf{L}^{\mathfrak{M}}/{=}$; the atoms below $\tilde{a}$ are the classes $\widetilde{\mathrm{at}(x)}$ with $x \in S_a$ and $0 < \|x\| \leq 1$. Any upper bound $\tilde{b}$ of these atoms satisfies $S_b \supseteq S_a$, hence $\tilde{b} \geq \tilde{a}$, which makes $\tilde{a}$ their join.

For (CP), let $\tilde{a}$ be an element of $\mathbf{L}^{\mathfrak{M}}/{=}$, and let $\widetilde{\mathrm{at}(x)}$ be an atom not below $\tilde{a}$; then the subspace S_a does not contain x. By Proposition 6.2 and Lemma A.11, the subspace $S_{a \vee \mathrm{at}(x)}$ is $S_a + \mathbb{C}x$, in which S_a has codimension one. Since τ is an order embedding, no element lies strictly between $\tilde{a}$ and $(a \vee \mathrm{at}(x))^{\sim}$. Condition (I) holds because the join of two distinct atoms corresponds to a plane, which contains infinitely many rays, all in $\tau(\mathbf{L}^{\mathfrak{M}}/{=})$. The atoms $\widetilde{\mathrm{at}(e_n)}$, for $n \in \mathbb{N}$, give (D).

For (S), take $p_n = \widetilde{\mathrm{at}(e_n)}$ and $c_n = \mathrm{at}(\widetilde{\frac{1}{2}(e_n + e_{n+1})})$; since τ is injective and preserves the lattice operations and the orthocomplementation, it suffices to verify the condition for the images of these atoms. These images are the rays of e_n and of $e_n + e_{n+1}$; by (1), the harmonic conjugate of $\mathbb{C}(e_n + e_{n+1})$ with respect to $\mathbb{C}e_n$ and $\mathbb{C}e_{n+1}$ is $\mathbb{C}(e_n - e_{n+1})$, the orthogonal complement of $\mathbb{C}(e_n + e_{n+1})$ in their plane, because $\|e_n\| = \|e_{n+1}\|$.

Each closed subspace S of H is the join in $L(H)$ of its rays, and it equals $\bigcap\{y^{\perp} : y \in S^{\perp},\ 0 < \|y\| \leq 1\}$, where each $y^{\perp} = S_{\mathrm{at}(y)^{\perp}}$ lies in the image of τ. These two facts make

the image of τ join-dense and meet-dense in $L(H)$. By Proposition A.5, the lattice $L(H)$ is therefore isomorphic to the MacNeille completion of $\mathbf{L}^{\mathfrak{M}}/{=}$.

Finally, let $A \subseteq \mathbf{L}^{\mathfrak{M}}/{=}$ have a join $\tilde{c}$ in $\mathbf{L}^{\mathfrak{M}}/{=}$, and let W be the closed span of the subspaces S_a, for $\tilde{a} \in A$. Since every element of A lies below $\tilde{c}$, the subspace W is contained in S_c. If y is a vector of $S_c \cap W^\perp$ with $0 < \|y\| \leq 1$, then every element of A lies below $(\mathrm{at}(y)^\perp)^\sim$, and hence so does $\tilde{c}$. The vector y, which lies in S_c, is then orthogonal to itself, which is absurd. Since $S_c = W + (S_c \cap W^\perp)$ by the projection theorem, it follows that $S_c = W$. □

Example 6.5. Let L_{fc} (for "finite dimension or finite codimension") be the set of finite-dimensional subspaces of the space ℓ^2 of square-summable complex sequences and their orthogonal complements. For finite-dimensional subspaces F and G of ℓ^2, the joins $F \vee G = F + G$, $F^\perp \vee G^\perp = (F \cap G)^\perp$ and $F \vee G^\perp = (F^\perp \cap G)^\perp$ in $L(\ell^2)$ belong to L_{fc}, and the De Morgan law places the meets in L_{fc} as well; the set L_{fc} is therefore a sub-ortholattice of $L(\ell^2)$. Take the closed unit ball of ℓ^2 as Hilbert sort and the standard unit vectors as the constants e_n, and take L_{fc}, with its lattice operations and orthocomplementation, as Solèr sort. Define $\pi(S, x) = P_S x$ and $\mathrm{at}(x) = \mathbb{C}x$, let $\mathrm{rep}(S)$ be a unit vector chosen in S for each $S \neq \{0\}$, and let $\mathrm{rep}(\{0\})$ be 0. Each truncation $R^{\pm}_{n,q}$ of a predicate of the Hilbert sort is interpreted by its intended value $\min\bigl(1, \max(0, \pm n(R - q))\bigr)$. Axioms (Sub1) and (Sub2) hold by Theorem 5.1, and the other axioms of T_0 by construction; the result is therefore a standard model of T_0, in which $[S = S'] = 1$ exactly when $S = S'$. Its Solèr sort L_{fc}, which is its own quotient, is not complete. Indeed, an upper bound in L_{fc} of the rays $\mathbb{C}e_{2n}$, for $n \in \mathbb{N}$, is infinite-dimensional and hence of the form $G^\perp$ for a finite-dimensional subspace G orthogonal to every e_{2n}. The closed span of the vectors e_{2n+1} is infinite-dimensional, whereas G is finite-dimensional; this span therefore contains a unit vector y orthogonal to G. The subspace $(G + \mathbb{C}y)^\perp$, which contains every e_{2n} but not y, is an upper bound of the same rays that lies strictly below $G^\perp$. The rays $\mathbb{C}e_{2n}$ therefore have no join in L_{fc}.

We now show that Gleason's representation extends to the quotient of the Solèr sort in every standard model of T_0. Its σ-additive states correspond to density operators on the model's Hilbert space H, which represent quantum states. Because the quotient may be incomplete, countable additivity is required only for orthogonal families whose join exists there.

A *density operator* on H is a positive operator ρ whose trace $\operatorname{Tr}\rho = \sum_k \langle \rho u_k \mid u_k \rangle$, computed in any orthonormal basis $(u_k)_{k\in\mathbb{N}}$ of H, equals 1. A *σ-additive state* on an ortholattice Λ is a map $\nu\colon \Lambda \to [0,1]$ with $\nu(1) = 1$ such that $\nu(\bigvee_{n\in J} a_n) = \sum_{n\in J} \nu(a_n)$ for every family $(a_n)_{n\in J}$ of mutually orthogonal elements of Λ, indexed by a set $J \subseteq \mathbb{N}$, whose join exists in Λ.

Theorem 6.6 (Gleason's theorem for T_0). *Let $\mathfrak{M}$ be a standard model of T_0, let H be the Hilbert space whose closed unit ball is its Hilbert sort, and, for each element a of its Solèr sort $\mathbf{L}^{\mathfrak{M}}$, let P_a be the orthogonal projection onto the closed subspace spanned by the members of a. The σ-additive states on $\mathbf{L}^{\mathfrak{M}}/{=}$ are exactly the maps $\tilde{a} \mapsto \operatorname{Tr}(\rho P_a)$, where ρ ranges over the density operators on H, and each such state determines its density operator.*

Proof. A density operator gives a σ-additive state by taking traces of the projections; conversely, the values of a σ-additive state on rays form a frame function, to which Gleason's theorem applies. Let ρ be a density operator on H. The map $\tilde{a} \mapsto \operatorname{Tr}(\rho P_a)$ is well defined, since P_a depends only on $\tilde{a}$. Suppose that the elements $\tilde{a}_n$, for n in a set $J \subseteq \mathbb{N}$, are mutually orthogonal and have a join $\tilde{c}$ in $\mathbf{L}^{\mathfrak{M}}/{=}$. By Proposition 6.2 and Theorem 6.4, the

subspace S_c is then the closure of the sum of the mutually orthogonal subspaces S_{a_n}, whence $P_c x = \sum_{n\in J} P_{a_n} x$ for every $x \in H$. We compute the traces in an orthonormal basis of H assembled from orthonormal bases of the subspaces S_{a_n} and of $S_c^\perp$; since the terms $\langle \rho u \mid u\rangle$ are nonnegative, they may be grouped by n, and we obtain $\mathrm{Tr}(\rho P_c) = \sum_{n\in J} \mathrm{Tr}(\rho P_{a_n})$. Moreover, the map $\tilde{a} \mapsto \mathrm{Tr}(\rho P_a)$ takes the value $\mathrm{Tr}\,\rho = 1$ at $\tilde{1}$, and its values lie in $[0,1]$.

Conversely, let ν be a σ-additive state on $\mathbf{L}^{\mathfrak{M}}/{=}$. Additivity on the orthogonal pair $\tilde{1}, \tilde{0}$, whose join is $\tilde{1}$, gives $\nu(\tilde{0}) = 0$. For a unit vector x, define $f(x) = \nu(\widetilde{\mathrm{at}(x)})$. For $S_a \neq \{0\}$, let $\mathcal{U}_a$ be an orthonormal basis of S_a, which is countable because H is separable. The atoms $\widetilde{\mathrm{at}(u)}$, for $u \in \mathcal{U}_a$, are mutually orthogonal, and their join in $\mathbf{L}^{\mathfrak{M}}/{=}$ is $\tilde{a}$. Indeed, the element $\tilde{a}$ is an upper bound of these atoms, and every upper bound $\tilde{b}$ of them satisfies $S_b \supseteq S_a$. Hence the value $\nu(\tilde{a})$ is the sum $\sum_{u\in\mathcal{U}_a} f(u)$.

Taking $a = 1$ shows that f is a nonnegative *frame function* of weight 1: its sum over every orthonormal basis of H is 1. Since H is separable of dimension at least three, Gleason's theorem [Gle57] gives a unique density operator ρ with $f(x) = \langle \rho x \mid x\rangle$ for every unit vector x. It follows that $\nu(\tilde{a}) = \sum_{u\in\mathcal{U}_a} \langle \rho u \mid u\rangle = \mathrm{Tr}(\rho P_a)$ whenever $S_a \neq \{0\}$; the same equality holds when $S_a = \{0\}$, since then $\tilde{a} = \tilde{0}$ and $\nu(\tilde{0}) = 0$. □

Historical note. Gleason proved his theorem in 1957 to answer a question of Mackey's. Mackey had asked whether every countably additive assignment of probabilities to the closed subspaces of a separable Hilbert space is given by the Born rule [Mac57, Gle57]. The Born rule gives a closed subspace, in the state described by a density operator ρ, the probability $\mathrm{Tr}(\rho P)$, where P is the orthogonal projection onto the subspace. Frame functions are the heart of Gleason's proof: he shows first, by means of spherical harmonics, that the continuous frame functions on the unit sphere of $\mathbb{R}^3$ come from quadratic forms. Gleason then shows, in the hardest step, that every nonnegative frame function on this sphere is continuous. Finally, he passes from $\mathbb{R}^3$ to real and complex separable Hilbert spaces of dimension at least three. Gleason noted that the restriction on the dimension is essential: in dimension two the theorem fails. Cooke, Keane and Moran gave in 1985 an elementary proof, longer than Gleason's but with fewer prerequisites [CKM85].

7. Standard models of T_1 and T: Gleason's theorem, the range of the measure and definable sets of values

In the standard model of T_0 in Example 6.5, the Solèr sort represents only the subspaces of finite dimension or finite codimension, so T_0 does not require every closed subspace to be represented. Every closed subspace of a separable Hilbert space is, however, spanned by a square-summable sequence of vectors, and the theory T_1 adds the sequence sort and the span function span, with their axioms, to supply these sequences and their spans. In a standard model of T_1 we first identify the sequence sort with the closed unit ball of the space of square-summable sequences of vectors of the Hilbert space, and then show that span represents the closed spans of these sequences; the quotient of the Solèr sort is consequently isomorphic to the whole lattice of closed subspaces of the Hilbert space, the representation needed for Theorem B and Proposition 7.5. In Section 7.1 we use the range of the measure to show that the sets of probabilities defined uniformly in standard models are exactly the σ-projective subsets of $[0,1]$.

Until Section 7.1, fix a standard model $\mathfrak{M}$ of T_1. Let H and K be the Hilbert spaces whose closed unit balls are its Hilbert and sequence sorts, and write $\mathbf{L}^{\mathfrak{M}}$ for its Solèr sort. The reduct of $\mathfrak{M}$ to the vocabulary of T_0 is a standard model of T_0, so the geometric representation of

Section 6 remains available. We retain its notation: P_a extends $\pi(a,\cdot)$, S_a is its range, $\tilde{a}$ is the class of a in $\mathbf{L}^{\mathfrak{M}}/{=}$, and τ sends $\tilde{a}$ to S_a. In the intended interpretation, the map κ_n reads the entry in position n of a sequence, and the map ι_n places a vector in that position (see Section 4); the first lemma identifies the sequence sort accordingly.

Lemma 7.1. *For each $n \in \mathbb{N}$, the map κ_n from the sequence sort to the Hilbert sort extends to a bounded linear map from K to H, and the map $s \mapsto (\kappa_n s)_{n\in\mathbb{N}}$ is a linear isometry of K onto the space $\ell^2(\mathbb{N}, H)$ of sequences $(x_n)_{n\in\mathbb{N}}$ in H with $\sum_n \|x_n\|^2 < \infty$.*

Proof. By (K2), each map κ_n is linear on the unit ball of K, and the maps ι_n extend to linear isometries from H into K with mutually orthogonal ranges. By the same axiom, we have $\langle \kappa_m s \mid x\rangle = \langle s \mid \iota_m x\rangle$. It follows that $\|\kappa_m s\| \le \|s\|$, and $\kappa_m \iota_n x = \delta_{mn} x$. Each map κ_m therefore extends to a bounded linear map from K to H.

For s and t in the unit ball of K, axiom (K3) gives $\langle \mathrm{res}_N s \mid t\rangle = \sum_{n<N}\langle \kappa_n s \mid \kappa_n t\rangle = \sum_{n<N}\langle \iota_n\kappa_n s \mid t\rangle$, whence $\mathrm{res}_N s = \sum_{n<N} \iota_n\kappa_n s$. Since the identity $\langle \kappa_m s \mid x\rangle = \langle s \mid \iota_m x\rangle$ makes κ_m the adjoint of the isometry ι_m, the operator $\sum_{n<N} \iota_n\kappa_n$ is the orthogonal projection onto the sum of the mutually orthogonal ranges of $\iota_0, \dots, \iota_{N-1}$. These projections increase with N, whence the norms $\|s - \mathrm{res}_N s\|$ decrease, and by the second part of (K3) their infimum, and hence their limit, is 0. Consequently, the numbers $\langle \mathrm{res}_N s \mid s\rangle = \sum_{n<N} \|\kappa_n s\|^2$ converge to $\|s\|^2$, which is therefore equal to $\sum_n \|\kappa_n s\|^2$, and the map $s \mapsto (\kappa_n s)_{n\in\mathbb{N}}$ is isometric on the unit ball and, being linear, on K.

Given $(x_n)_{n\in\mathbb{N}} \in \ell^2(\mathbb{N}, H)$, the partial sums $\sum_{n<N} \iota_n x_n$ form a Cauchy sequence in K, since the maps ι_n are isometries with mutually orthogonal ranges. Since K is complete, these partial sums converge to some s, which satisfies $\kappa_n s = x_n$ by the continuity of κ_n. □

Lemma 7.2. *For every s in the unit ball of K, the subspace $S_{\mathrm{span}(s)}$ is the closed span of the entries $\kappa_n s$.*

Proof. Let W be this closed span. By (Sp1), the subspace $S_{\mathrm{span}(s)}$ contains every entry of s, and hence it contains W. Conversely, let y be a vector of $W^\perp$ with $0 < \|y\| \le 1$, and let $b = \mathrm{at}(y)^\perp$; by Lemma 5.4, the subspace S_b is $y^\perp$. Every entry $\kappa_n s$ lies in S_b, and so does the subspace $S_{\mathrm{at}(\kappa_n s)}$, which is $\mathbb{C}\kappa_n s$ or $\{0\}$; by Lemma 5.4, we therefore have $[\mathrm{at}(\kappa_n s) \le b] = 1$ for every n. Axiom (Sp2) gives $[\mathrm{span}(s) \le b] = 1$, that is, $S_{\mathrm{span}(s)} \subseteq y^\perp$. Since y was arbitrary, the subspace $S_{\mathrm{span}(s)}$ is orthogonal to $W^\perp$, whence $S_{\mathrm{span}(s)} \subseteq W^{\perp\perp} = W$. □

Theorem 7.3. *Let $\mathfrak{M}$ be a standard model of T_1, let H be the Hilbert space whose closed unit ball is its Hilbert sort, and, for each element a of its Solèr sort $\mathbf{L}^{\mathfrak{M}}$, let S_a be the closed subspace of H spanned by the members of a and let $\tilde{a}$ be the class of a in the quotient $\mathbf{L}^{\mathfrak{M}}/{=}$ of $\mathbf{L}^{\mathfrak{M}}$ by its equality, the relation $[a = b] = 1$ of Section 4. Then H is a separable infinite-dimensional complex Hilbert space, and the map $\tau\colon \tilde{a} \mapsto S_a$ is an isomorphism of ortholattices from $\mathbf{L}^{\mathfrak{M}}/{=}$ onto the lattice $L(H)$ of closed subspaces of H; in particular, this quotient is a Solèr lattice. Conversely, for every separable infinite-dimensional complex Hilbert space H', there is a standard model of T_1 whose Hilbert sort is the closed unit ball of H' and whose Solèr sort is $L(H')$.*

Proof. By Proposition 6.1, the space H is separable and infinite-dimensional. By Proposition 6.2, the map τ is an embedding; it remains to be shown that it is onto, for then the quotient $\mathbf{L}^{\mathfrak{M}}/{=}$, being isomorphic to $L(H)$, is a Solèr lattice by Theorem A.18. Let S be a closed subspace of H, and let $(u_n)_{n\in\mathbb{N}}$ enumerate an orthonormal basis of S, which is countable by Proposition 6.1, followed by zeros if the basis is finite. By Lemma 7.1 there is s in the

unit ball of K with $\kappa_n s = 2^{-n-1}u_n$ for all n, since $\sum_n 4^{-n-1} < 1$, and the subspace $S_{\mathrm{span}(s)}$ is S by Lemma 7.2.

For the converse, we build the model sort by sort. The Hilbert sort is the closed unit ball of H', and the constants e_n are interpreted as an orthonormal basis $(e_n)_{n\in\mathbb{N}}$ of H'. The Solèr sort is $L(H')$, with $\pi(S,x) = P_S x$ for the orthogonal projection P_S onto S and with $\mathrm{at}(x) = \mathbb{C}x$. For each $S \neq \{0\}$ we let $\mathrm{rep}(S)$ be a unit vector chosen in S, and we let $\mathrm{rep}(\{0\})$ be 0. The sequence sort is the closed unit ball of $\ell^2(\mathbb{N}, H')$, on which the function ι_n places a vector x in position n of a sequence whose other entries are 0, and the function κ_n reads the entry of a sequence s in position n. The function res_N keeps the entries of s in the positions $n < N$ and replaces the others by 0, while the constant ε is the sequence $(2^{-n-1}e_n)_{n\in\mathbb{N}}$, whose norm is less than 1. The function span sends each element s of the sequence sort to the closed span of its entries. Finally, each truncation $R^{\pm}_{n,q}$ is interpreted by its intended value $\min\big(1, \max(0, \pm n(R-q))\big)$, so that the axioms (Tr) hold. Axioms (Sub1) and (Sub2) hold by Theorem 5.1, and the others by construction; in this model, the formula $[S = S']$ takes the value 1 exactly when $S = S'$. □

Proof of Theorem B. The quotient of the Solèr sort has now been identified with $L(H)$; it remains to identify the measure. Axiom (M4) gives finite additivity, and (M5) will turn the measure of a span into a sum over an orthonormal basis. Together these yield the σ-additive state needed for Theorem 6.6. Let $\mathfrak{M}$ be a standard model of T. Its reduct to the vocabulary of T_1 is a standard model of T_1, to which the notation fixed at the beginning of this section applies. By Theorem 7.3, the space H is a separable infinite-dimensional complex Hilbert space, and the map τ is an isomorphism of $\mathbf{L}^{\mathfrak{M}}/{=}$ onto $L(H)$. By (M3), the value $\mu(a)$ depends only on the class $\tilde{a}$, and we write $\mu(\tilde{a})$ for it.

Step 1: finite additivity. Let a and c be elements of $\mathbf{L}^{\mathfrak{M}}$ with $S_c \subseteq S_a^{\perp}$. By Lemma 5.4 and Corollary 5.5, the elements $c \wedge a^{\perp}$ and c have the same class, and so have $a \vee (c \wedge a^{\perp})$ and $a \vee c$. Axioms (M3) and (M4) therefore give $\mu(a \vee c) = \mu(a) + \mu(c)$. Taking $a = c = 0$ gives $\mu(0 \vee 0) = 2\mu(0)$, and since $0 \vee 0$ and 0 have the same class, (M3) turns this equation into $\mu(0) = 2\mu(0)$, whence $\mu(0) = 0$.

Step 2: sums along orthonormal bases. Let a be an element of $\mathbf{L}^{\mathfrak{M}}$, and let $(u_k)_{k\in\mathbb{N}}$ enumerate an orthonormal basis of S_a, followed by zeros if it is finite. As in the proof of Theorem 7.3, Lemma 7.1 gives an element s of the unit ball of K with $\kappa_k s = 2^{-k-1}u_k$ for every k, and by Lemma 7.2 the element $\mathrm{span}(s)$ has the class $\tilde{a}$. For $u_k \neq 0$, the subspaces $S_{\mathrm{at}(\kappa_k s)}$ and $S_{\mathrm{at}(u_k)}$ are both $\mathbb{C}u_k$ by Proposition 6.3, and for $u_k = 0$ both are $\{0\}$ by part (4) of Lemma 5.4. By Step 1 and (M3), the measure of $c_N(s)$ is therefore $\sum_{k<N} \mu(\mathrm{at}(u_k))$. These partial sums increase with N, since their terms are nonnegative by (M1), and axioms (M5) and (M3) therefore give

$$\mu(a) = \sum_k \mu(\mathrm{at}(u_k)). \qquad (*)$$

Step 3: application of Gleason's theorem. We show that the map $\tilde{a} \mapsto \mu(\tilde{a})$ is a σ-additive state on $\mathbf{L}^{\mathfrak{M}}/{=}$, in order to apply Theorem 6.6. Its values lie in $[0,1]$ by (M1), and by (M2) it takes the value 1 at $\tilde{1}$. Let $(\tilde{a}_n)_{n\in J}$, for a set $J \subseteq \mathbb{N}$, be a family of mutually orthogonal elements of $\mathbf{L}^{\mathfrak{M}}/{=}$. This family has a join $\tilde{c}$, because the map τ is onto $L(H)$. By Proposition 6.2 and Theorem 6.4, the subspaces S_{a_n} are mutually orthogonal, and the closure of their sum is S_c; orthonormal bases of the subspaces S_{a_n} therefore form together an orthonormal basis of S_c. Applied to c, the identity $(*)$ expresses $\mu(c)$ as the sum of the values $\mu(\mathrm{at}(u))$ over the vectors u of this basis. Applied to each a_n, the identity expresses $\mu(a_n)$ as the sum of the same values over the vectors that come from S_{a_n}. Since these values are

nonnegative by (M1), the first sum may be grouped by n, whence $\mu(\tilde{c}) = \sum_{n\in J} \mu(\tilde{a}_n)$. By Theorem 6.6, applied to the reduct of $\mathfrak{M}$ to the vocabulary of T_0, there is a unique density operator ρ on H with $\mu(a) = \mathrm{Tr}(\rho P_a)$ for every $a \in \mathbf{L}^{\mathfrak{M}}$.

The converse. Let a separable infinite-dimensional complex Hilbert space H' and a density operator ρ on H' be given. We expand a standard model of T_1 built for H' as in the proof of Theorem 7.3 by $\mu(a) = \mathrm{Tr}(\rho P_a)$, where P_a is the orthogonal projection onto the closed subspace a of H', and we interpret the truncations of μ by their intended values, so that the axioms (Tr) for μ hold. Axiom (M3) holds because the formula $[a = b]$ takes the value 1 only when $a = b$. Axioms (M1) and (M2) hold because every orthogonal projection P satisfies $0 \le \mathrm{Tr}(\rho P) \le 1$ and because the trace of ρ is 1. Axiom (M4) holds because the subspaces a and $b \wedge a^{\perp}$ are orthogonal, whence the projection onto their join is the sum of their projections. For (M5), let s be an element of the sequence sort. The subspaces $c_N(s)$ increase with N, and the closure of their union is the closed span $\mathrm{span}(s)$ of the entries of s. The projections onto them therefore converge strongly to $P_{\mathrm{span}(s)}$, and since ρ is of trace class, the numbers $\mathrm{Tr}(\rho P_{c_N(s)})$ increase to $\mathrm{Tr}(\rho P_{\mathrm{span}(s)})$. □

The representation in Theorem B expresses the measure through a density operator, but Proposition 7.5 also needs its full range. For a density operator ρ that is not a rank-one projection, every ray probability is at most $\|\rho\| < 1$, so the following lemma considers all closed subspaces, not only rays.

Lemma 7.4. *For every separable infinite-dimensional complex Hilbert space H and every density operator ρ on H, the set $\{\mathrm{Tr}(\rho P_S) : S \in L(H)\}$ is the interval $[0, 1]$.*

Proof. The values 0 and 1 are attained at $\{0\}$ and H, respectively. Fix α with $0 < \alpha < 1$. We start with a subspace spanned by eigenvectors whose probability reaches or exceeds α, and rotate its last basis vector until the probability falls below α. Continuity will then give the desired value. Since ρ is a positive compact operator of trace 1, there are orthonormal eigenvectors $v_1, v_2, \ldots$ (finitely or infinitely many) whose eigenvalues $\eta_1 \ge \eta_2 \ge \cdots > 0$ are the nonzero eigenvalues of ρ, repeated according to multiplicity, and these eigenvalues satisfy $\sum_i \eta_i = 1$. Write $\eta_{\le k} = \eta_1 + \cdots + \eta_k$, with $\eta_{\le 0} = 0$, and choose an integer $k \ge 0$ with $\eta_{\le k} < \alpha \le \eta_{\le k+1}$.

On the unit vectors of $W = \{v_1, \ldots, v_{k+1}\}^{\perp}$, the quantity $\langle \rho x \mid x\rangle$ has infimum 0. Indeed, if ρ has only finitely many nonzero eigenvalues, then its kernel has finite codimension in the infinite-dimensional space H, and it therefore meets W, which also has finite codimension, in a nonzero vector; otherwise the eigenvectors v_i with $i > k+1$ lie in W, and their eigenvalues η_i tend to 0. Choose a unit vector $u \in W$ with $\langle \rho u \mid u\rangle < \alpha - \eta_{\le k}$.

For $0 \le \vartheta \le 1$, let w_ϑ be the unit vector $\sqrt{1-\vartheta}\, v_{k+1} + \sqrt{\vartheta}\, u$, and let S_ϑ be the subspace spanned by $v_1, \ldots, v_k$ and w_ϑ. Because u is orthogonal to $v_1, \ldots, v_{k+1}$, the vectors $v_1, \ldots, v_k, w_\vartheta$ form an orthonormal basis of S_ϑ, and the trace of ρP_{S_ϑ} is the sum of the values of $\langle \rho x \mid x\rangle$ at these vectors. In the expansion of $\langle \rho w_\vartheta \mid w_\vartheta\rangle$, the cross terms are multiples of $\langle \rho v_{k+1} \mid u\rangle = \eta_{k+1}\langle v_{k+1} \mid u\rangle$ and of its conjugate, and they therefore vanish. Consequently, we have

$$\mathrm{Tr}(\rho P_{S_\vartheta}) = \eta_{\le k} + (1-\vartheta)\eta_{k+1} + \vartheta\langle \rho u \mid u\rangle,$$

a continuous function of ϑ that equals $\eta_{\le k+1} \ge \alpha$ at $\vartheta = 0$ and is smaller than α at $\vartheta = 1$. By the intermediate value theorem, it takes the value α. □

Proposition 7.5 (The range of the measure). *In every standard model of T, the range of μ is the whole interval $[0, 1]$. Consequently, for every Borel set $B \subseteq \mathbb{R}$ that meets $[0, 1]$,*

the theory T, together with the axiom $\beta_B(a) \leq -1$, which says that no value of μ lies in B (with β_B as in Lemma 4.1), has no standard models.

Proof. Let $\mathfrak{M}$ be a standard model of T, and let H be the Hilbert space whose closed unit ball is its Hilbert sort. By Theorem B, there is a density operator ρ on H with $\mu(a) = \mathrm{Tr}(\rho P_a)$ for every element a of the Solèr sort, and the map $\tilde{a} \mapsto S_a$ is onto $L(H)$. By Lemma 7.4, the range of μ is $[0,1]$. Let $B \subseteq \mathbb{R}$ be a Borel set that meets $[0,1]$. Some element a of the Solèr sort then satisfies $\mu(a) \in B$, and the formula $\beta_B(a)$ takes the value 1 at this element. Hence $\mathfrak{M}$ does not satisfy the axiom $\beta_B(a) \leq -1$, which says that no value of μ lies in B. □

⧖ **Historical note.** Sierpiński proved in 1922 that a measure without atoms takes, on the measurable subsets of a set, every value between 0 and the measure of the set [Sie22]. A $\{0,1\}$-valued measure, which is what an assignment of truth values amounts to, is the opposite extreme: the whole space is an atom. By contrast, the measures $S \mapsto \mathrm{Tr}(\rho P_S)$ that density operators ρ define on the closed subspaces of a separable infinite-dimensional Hilbert space take every value in $[0,1]$ (see Lemma 7.4). Such a measure takes all these values even when ρ is a pure state, the orthogonal projection onto a ray, whose classical counterpart would be a point mass.

7.1. **Definable sets of values.** The Borel tests of Lemma 4.1 determine whether a probability belongs to a prescribed Borel set, using connectives alone. For formulas depending on an additional lattice variable only through its measure, quantification over the variable performs a projection along $[0,1]$, since every value in this interval is a value of μ in every standard model of T (see Proposition 7.5). We show in Theorem 7.7 that the sets obtained uniformly in standard models are exactly the σ-projective subsets of the interval. The exclusion result of Proposition 7.5 therefore extends from nonempty Borel sets of probabilities to all nonempty sets in this class.

To compare the operations on formulas with operations on sets, we use the least classes containing the Borel sets and closed under complements, countable unions and projections along $\mathbb{N}^{\mathbb{N}}$. Aguilera, Müller and Schlicht [AMS21] study the resulting σ-projective sets on finite powers of $\mathbb{N}^{\mathbb{N}}$. The proof of Theorem 7.7 below also uses cubes of probabilities and a space of orthogonal projections, so we define the classes on Polish spaces and establish closure under Borel preimages and projections along Polish spaces. Let Γ be the least assignment of a class $\Gamma(X)$ of subsets of X to each Polish space X with the following two properties: each class $\Gamma(X)$ contains the Borel subsets of X and is closed under complements in X and countable unions; and, whenever $A \in \Gamma(X \times \mathbb{N}^{\mathbb{N}})$, the set $\{x \in X : (x,z) \in A \text{ for some } z \in \mathbb{N}^{\mathbb{N}}\}$ belongs to $\Gamma(X)$. A subset of a Polish space X is *σ-projective* when it belongs to $\Gamma(X)$. In particular, every projective set [Kec95, Section 37] is σ-projective, and so is every countable union of projective sets.

Definition 7.6. A set $D \subseteq [0,1]$ is *definable over T* if there is a formula $\gamma(a)$ of real-valued $L_{\omega_1\omega}$, with the single free variable a of the Solèr sort, whose value, in every standard model of T and at every element a of its Solèr sort, is 1 if $\mu(a) \in D$ and -1 otherwise; we say that *γ defines D*.

Only subsets of $[0,1]$ matter here, because (M1) confines μ to this interval. A formula that tests membership in $E \subseteq \mathbb{R}$ in every standard model therefore defines $E \cap [0,1]$ in the sense of Definition 7.6.

Theorem 7.7. *A subset of $[0,1]$ is definable over T if and only if it is σ-projective.*

Consequently, if $D \subseteq [0,1]$ is a nonempty σ-projective set and $\gamma(a)$ defines D, then the theory T together with the axiom $\gamma(a) \leq -1$, which in a standard model says that no value of μ lies in D, has no standard models.

To define the sets, we encode membership in tuples of measure values and use the full range of μ to express projection along $[0,1]$ by quantification over the Solèr sort. For the converse, a formula that works in every standard model must work in one whose symbols are interpreted by Borel functions on Polish spaces. In this model we follow the descriptive complexity of the formula's values through the connectives and quantifiers, then pull back along a continuous map assigning to each probability $t \in [0,1]$ a projection of measure t.

Closure properties used in the proof. On the finite powers of $\mathbb{N}^\mathbb{N}$, the σ-projective sets are the sets of Aguilera, Müller and Schlicht. Their class is the least that contains the open sets and is closed under complements, countable unions and projections along $\mathbb{N}^\mathbb{N}$. This class is therefore contained in Γ, which has these properties on the finite powers of $\mathbb{N}^\mathbb{N}$. Conversely, for a Polish space X, consider the sets $A \subseteq X$ such that $g^{-1}[A]$ belongs to the class of Aguilera, Müller and Schlicht for every $k \geq 1$ and every Borel map $g\colon (\mathbb{N}^\mathbb{N})^k \to X$. The class of Aguilera, Müller and Schlicht contains the Borel sets and is closed under the operations that define Γ. Hence these sets A include the Borel sets and are closed under complements, countable unions and projections along $\mathbb{N}^\mathbb{N}$. By the minimality of Γ, they include $\Gamma(X)$. For $X = (\mathbb{N}^\mathbb{N})^k$, taking for g the identity map then shows that every set of $\Gamma(X)$ belongs to the class of Aguilera, Müller and Schlicht.

The σ-projective sets are closed under preimages by Borel maps between Polish spaces. Indeed, the sets $A \subseteq X$ such that $g^{-1}[A]$ is σ-projective for every Borel map g from a Polish space into X include the Borel sets and are closed under complements and countable unions. They are also closed under projections along $\mathbb{N}^\mathbb{N}$, since, for a set $A \subseteq X \times \mathbb{N}^\mathbb{N}$ and a Borel map g from a Polish space Z into X, the preimage under g of the projection of A is the projection of the preimage of A under the Borel map $g \times \mathrm{id}$ from $Z \times \mathbb{N}^\mathbb{N}$ into $X \times \mathbb{N}^\mathbb{N}$. These sets therefore include every σ-projective set.

The σ-projective sets are also closed under projections along every nonempty Polish space Y. Given a continuous surjection $p\colon \mathbb{N}^\mathbb{N} \to Y$ [Kec95, Theorem 7.9] and a σ-projective set $A \subseteq X \times Y$, the projection of A to X is the projection along $\mathbb{N}^\mathbb{N}$ of the preimage of A under $\mathrm{id} \times p$, and it is therefore σ-projective.

A function f from a Polish space to $\mathbb{R}$ is Γ-*measurable* when the set $\{f > r\}$ is σ-projective for every real number r. The set $\{f \geq r\}$, which is the intersection of the sets $\{f > r - \frac{1}{n}\}$, is then σ-projective as well.

Proof of Theorem 7.7. Every σ-projective subset of $[0,1]$ is definable. For $m \geq 1$, consider a formula $\gamma(a_1, \ldots, a_m)$ with free variables $a_1, \ldots, a_m$ of the Solèr sort. The formula γ *strongly defines* a set $A \subseteq [0,1]^m$ if it has a priori bound at most 3, is two-valued in every model of T and, in every standard model of T, equals 1 at $(a_1, \ldots, a_m)$ exactly when $(\mu(a_1), \ldots, \mu(a_m)) \in A$. Let $\mathcal{A}$ be the class of sets $A \subseteq [0,1]^m$, for $m \geq 1$, that some formula strongly defines. We show that $\mathcal{A}$ contains every σ-projective subset of every cube $[0,1]^m$; a formula that strongly defines a set $D \subseteq [0,1]$ defines it in the sense of Definition 7.6.

First, $\mathcal{A}$ contains the Borel sets. We repeat the proof of Lemma 4.1 with the closed boxes $[q_1, q_1'] \times \cdots \times [q_m, q_m']$, for rational numbers $q_i \leq q_i'$, in place of the closed intervals with rational endpoints. The formula for such a box is the minimum

$$\beta_{[q_1,q_1']}(a_1) \wedge \cdots \wedge \beta_{[q_m,q_m']}(a_m)$$

of the formulas of the original proof for the intervals $[q_i, q_i']$. For every Borel set $B \subseteq \mathbb{R}^m$, the modified proof gives a formula $\beta_B(a_1, \dots, a_m)$. In every structure for $\mathcal{L}$ that satisfies (Tr), its value at every tuple of elements of the Solèr sort is 1 if $(\mu(a_1), \dots, \mu(a_m)) \in B$ and -1 otherwise. The formula $\beta_B(a_1, \dots, a_m)$ strongly defines $B \cap [0,1]^m$. Next, if γ strongly defines $A \subseteq [0,1]^m$, then the formula $-\gamma$ strongly defines $[0,1]^m \setminus A$. If γ_n strongly defines $A_n \subseteq [0,1]^m$ for every $n \in \mathbb{N}$, then the join $\bigvee_n \gamma_n$, whose members have a priori bounds at most 3, strongly defines $\bigcup_n A_n$. Finally, if $\gamma(a_1, \dots, a_m, b)$ strongly defines $A \subseteq [0,1]^{m+1}$, then the formula $\sup_b \gamma(a_1, \dots, a_m, b)$ is two-valued in every model of T. Indeed, the Solèr sort contains the constant 0, and a supremum of values in $\{-1, 1\}$ over a nonempty set is attained. In a standard model of T, the formula $\sup_b \gamma(a_1, \dots, a_m, b)$ takes the value 1 at $(a_1, \dots, a_m)$ exactly when $(\mu(a_1), \dots, \mu(a_m), \mu(b)) \in A$ for some element b. By Proposition 7.5, such an element b exists exactly when $(\mu(a_1), \dots, \mu(a_m), t) \in A$ for some $t \in [0,1]$. Thus $\mathcal{A}$ contains the Borel subsets of the cubes $[0,1]^m$ and is closed under complements, countable unions and projections along $[0,1]$.

The formulas above implement projections along $[0,1]$, whereas Γ was defined using projections along $\mathbb{N}^{\mathbb{N}}$. A Borel isomorphism between these spaces lets us compare the two constructions. Consider the least class Γ_0 of subsets of the cubes $[0,1]^m$, for $m \geq 1$, containing their Borel subsets and closed under complements in the cube, countable unions and projections from $[0,1]^{m+1}$ to $[0,1]^m$. We have already shown that the class Γ_0 is contained in $\mathcal{A}$.

Fix a Borel isomorphism χ of $[0,1]$ onto $\mathbb{N}^{\mathbb{N}}$ [Kec95, Theorem 15.6]. For each Polish space X, let $\Gamma'(X)$ be the class of sets $A \subseteq X$ such that $g^{-1}[A] \in \Gamma_0$ for every $m \geq 1$ and every Borel map $g\colon [0,1]^m \to X$. Each class $\Gamma'(X)$ contains the Borel subsets of X and is closed under complements in X and countable unions. Let $A \in \Gamma'(X \times \mathbb{N}^{\mathbb{N}})$ be given, and let $g\colon [0,1]^m \to X$ be a Borel map. Since χ is onto $\mathbb{N}^{\mathbb{N}}$, we have

$$g^{-1}\big[\{x : \exists z \in \mathbb{N}^{\mathbb{N}}\ (x,z) \in A\}\big] = \{u \in [0,1]^m : \exists t \in [0,1]\ (u,t) \in (g \times \chi)^{-1}[A]\},$$

the projection along $[0,1]$ of a set of Γ_0, because $g \times \chi$ is a Borel map on $[0,1]^{m+1}$; this projection belongs to Γ_0. Hence the assignment Γ' is closed under projections along $\mathbb{N}^{\mathbb{N}}$, and we have $\Gamma(X) \subseteq \Gamma'(X)$ for every Polish space X. For a σ-projective set $A \subseteq [0,1]^m$, taking for g the identity shows that $A \in \Gamma_0$ and therefore that $A \in \mathcal{A}$.

Every definable subset of $[0,1]$ *is* σ*-projective.* Let $\gamma(a)$ define a set $D \subseteq [0,1]$. Let $H = \ell^2(\mathbb{N})$, let v and w be orthogonal unit vectors of H, and let ρ be the projection onto $\mathbb{C}v$, a density operator. By Theorem B(2) and its proof, there is a standard model $\mathfrak{M}$ of T that can be described as follows. Its Hilbert sort is the closed unit ball B_H of H, its sequence sort is the closed unit ball B_K of $K = \ell^2(\mathbb{N}, H)$, and its Solèr sort is $L(H)$. The symbols are interpreted as in the proof of Theorem B, with the standard unit vectors of $\ell^2(\mathbb{N})$ as the constants e_n and with $\mu(S) = \operatorname{Tr}(\rho P_S)$. There, the value $\operatorname{rep}(S)$ may be any unit vector of S when $S \neq \{0\}$, and here we let it be $P_S e_n / \|P_S e_n\|$ for the least n with $P_S e_n \neq 0$. We identify $L(H)$ with the set of orthogonal projections on H, with the strong operator topology. Fix a dense sequence $(x_k)_{k\in\mathbb{N}}$ in B_H, and give the set of operators of norm at most 1 the strong operator topology. For such operators, convergence at the vectors x_k implies convergence at every vector. The map $Q \mapsto (Qx_k)_{k\in\mathbb{N}}$ is therefore a homeomorphism of the set of operators of norm at most 1 onto its image in the Polish space $B_H^{\mathbb{N}}$. The image is closed: if $(y_k)_{k\in\mathbb{N}}$ is the limit of sequences $(Q_j x_k)_{k\in\mathbb{N}}$ in the image, then we have $\|\sum_k c_k y_k\| \leq \|\sum_k c_k x_k\|$ for every finitely supported family of complex numbers c_k. The map $x_k \mapsto y_k$ therefore extends to an operator of norm at most 1. The projections form a closed subset, since a strong limit Q of projections satisfies $\langle Qx \mid x\rangle = \|Qx\|^2 = \langle Q^*Qx \mid x\rangle$ for every x, whence

$Q = Q^*Q$ by polarization; the operator Q is therefore self-adjoint, and it is idempotent, since $Q^2 = Q^*Q = Q$. Thus the space $L(H)$ is Polish, and so are B_H and B_K with their norm topologies.

Every symbol of $\mathcal{L}$ is interpreted in $\mathfrak{M}$ by a Borel function on a product of these spaces. The vector operations, the inner products and the norms are continuous, as are the maps ι_n, κ_n and res_N, the constants and the orthocomplement $Q \mapsto I - Q$. The map $(Q, x) \mapsto Qx$, which interprets π, is continuous because $\|Qx - Q_0x_0\| \leq \|x - x_0\| + \|(Q - Q_0)x_0\|$. The measure, given by $\mu(Q) = \mathrm{Tr}(\rho Q) = \langle Qv \mid v\rangle$, is continuous as well.

The meet $Q \wedge Q'$ is the strong limit of $(QQ'Q)^n$ as $n \to \infty$, by Theorem 5.1 and the identity $(QQ'Q)^n = (QQ')^nQ$. As a pointwise limit of continuous maps, the meet is Borel, and the join is given by $Q \vee Q' = I - ((I - Q) \wedge (I - Q'))$. The truncations are continuous functions of the predicates that they truncate.

The map at sends $x \neq 0$ to the projection $y \mapsto \langle y \mid x\rangle x/\|x\|^2$, continuously, and 0 to 0. For rep, the set of projections Q with $Qe_k = 0$ for $k < n$ and $Qe_n \neq 0$ is Borel, and on it $\mathrm{rep}(Q) = Qe_n/\|Qe_n\|$ is continuous.

For span, the term $c_N(s) = \mathrm{at}(\kappa_0 s) \vee \cdots \vee \mathrm{at}(\kappa_{N-1}s)$ of axiom (M5) is interpreted by the projection onto the span of the entries $\kappa_0 s, \ldots, \kappa_{N-1}s$ of s, which is a Borel function of s, being a composition of the maps κ_n, at and $\vee$. These projections increase with N and converge strongly to the projection $\mathrm{span}(s)$ onto the closed span of the entries of s, and the map span, as a pointwise limit of Borel maps into the Polish space $L(H)$, is therefore Borel. It follows by induction on terms that every term is interpreted by a Borel function.

We show by induction on formulas that, for every formula $\varphi(\bar{x})$, the function $\bar{x} \mapsto \varphi^{\mathfrak{M}}(\bar{x})$ is Γ-measurable on the product of the Polish spaces that interpret the sorts of $\bar{x}$. Atomic formulas, those built from the truncations included, are continuous functions of Borel functions of terms, hence Borel. For the connectives $s \cdot \varphi$, $\varphi + \psi$ and $\varphi \vee \psi$, the set where the value exceeds r is the preimage of an open subset of $\mathbb{R}$ or $\mathbb{R}^2$ under the function given by the argument formula or pair of argument formulas. Write the open set as a countable union of rational open intervals in the unary case, or of rational open boxes in the binary cases. For rational endpoints $q < q'$, the preimage of (q, q') under $\varphi^{\mathfrak{M}}$ is $\{\varphi^{\mathfrak{M}} > q\} \setminus \{\varphi^{\mathfrak{M}} \geq q'\}$, and box preimages are intersections of the corresponding sets for $\varphi^{\mathfrak{M}}$ and $\psi^{\mathfrak{M}}$.

For a countable join, the set $\{\bigvee_n \varphi_n > r\}$ is the union $\bigcup_n \{\varphi_n > r\}$. For a supremum over a variable y, the set $\{\bar{x} : \sup_y \varphi(\bar{x}, y) > r\}$ is the projection of $\{(\bar{x}, y) : \varphi(\bar{x}, y) > r\}$ along the Polish space that interprets the sort of y. In each case the set is σ-projective.

Now the function $\gamma^{\mathfrak{M}}$ is Γ-measurable, and the set $\{Q \in L(H) : \gamma^{\mathfrak{M}}(Q) > 0\}$ is σ-projective. For $t \in [0, 1]$, let $\xi(t)$ be the projection $\mathrm{at}(\sqrt{t}\, v + \sqrt{1-t}\, w)$ onto the line spanned by the unit vector $\sqrt{t}\, v + \sqrt{1-t}\, w$. The map ξ is continuous and satisfies $\mu(\xi(t)) = |\langle v \mid \sqrt{t}\, v + \sqrt{1-t}\, w\rangle|^2 = t$. Since $\gamma^{\mathfrak{M}}$ takes the value 1 exactly at the elements Q with $\mu(Q) \in D$,

$$D = \{t \in [0, 1] : \gamma^{\mathfrak{M}}(\xi(t)) > 0\} = \xi^{-1}\big[\{\gamma^{\mathfrak{M}} > 0\}\big],$$

a continuous preimage of a σ-projective set, hence σ-projective.

The consequence stated after the theorem follows as in the proof of Proposition 7.5, with γ in place of β_B: in a standard model, some element a has $\mu(a) \in D$, and $\gamma(a)$ takes the value 1 there. □

The class of σ-projective subsets of $[0, 1]$ is strictly larger than that of projective ones. For each $n \geq 1$, the hierarchy theorem for projective sets [Kec95, Section 37] provides a $\mathbf{\Sigma}^1_{n+1}$ set $C_n \subseteq [0, 1]$ that is not $\mathbf{\Sigma}^1_n$, and we let h_n be a homeomorphism of $[0, 1]$ onto $[2^{-n}, 2^{-n+1}]$. The set $C = \bigcup_n h_n[C_n \cap (0, 1)]$ is a countable union of projective sets, and it contains none

of the points 2^{-m}. If C were $\mathbf{\Sigma}^1_k$, then for $n > k$ the set $h_n^{-1}[C]$, which equals $C_n \cap (0,1)$, would be $\mathbf{\Sigma}^1_k$. The set C_n, which differs from $C_n \cap (0,1)$ by at most two points, would then be $\mathbf{\Sigma}^1_k$ and hence $\mathbf{\Sigma}^1_n$, which is false. If C were $\mathbf{\Pi}^1_k$, it would be $\mathbf{\Sigma}^1_{k+1}$, which is excluded in the same way. Hence C is not projective, and by Theorem 7.7 it is definable over T.

The upper-bound argument also relates the complexity of a defined set to the nesting of quantifiers in its defining formula, and shows why defining the nonprojective set C above requires a countable join whose members have unbounded finite quantifier depths. We define the *quantifier depth* $\mathrm{qd}(\varphi)$ of a formula φ, a countable ordinal, by induction. Atomic formulas have depth 0, and the depth of $s \cdot \varphi$, of $\varphi + \psi$ and of $\varphi \vee \psi$ is the largest depth of the formulas to which the connective is applied. The depth of a countable join is given by $\mathrm{qd}(\bigvee_n \varphi_n) = \sup_n \mathrm{qd}(\varphi_n)$, and that of a supremum by $\mathrm{qd}(\sup_x \varphi) = \mathrm{qd}(\varphi) + 1$.

Proposition 7.8 (Quantifier depth and projective complexity). *Let $\mathfrak{M}$ be the standard model of T constructed in the proof of Theorem 7.7, whose symbols are interpreted by Borel functions, and let $k < \omega$.*

1. *For every formula φ with $\mathrm{qd}(\varphi) \leq k$ and every real number r, the set $\{\varphi^{\mathfrak{M}} > r\}$ is $\mathbf{\Delta}^1_{k+1}$.*
2. *Every set $D \subseteq [0,1]$ definable over T by a formula of quantifier depth at most k is $\mathbf{\Delta}^1_{k+1}$.*
3. *Every formula that defines the set C above over T has quantifier depth at least ω, and contains a countable join whose members have unbounded finite quantifier depths.*

Proof. For part (1), the induction on formulas given in the proof of Theorem 7.7 applies. For $k = 0$, the formula φ contains no supremum over a variable, and these sets are Borel. For $k \geq 1$, if $\mathrm{qd}(\sup_y \varphi) \leq k$, then the formula φ has depth at most $k-1$, and the set $\{(\sup_y \varphi)^{\mathfrak{M}} > r\}$ is the projection of a $\mathbf{\Delta}^1_k$ set. This projection is therefore $\mathbf{\Sigma}^1_k$, and hence $\mathbf{\Delta}^1_{k+1}$. The other cases use only countable unions, countable intersections and complements, under which $\mathbf{\Delta}^1_{k+1}$ is closed. For part (2), if a formula γ with $\mathrm{qd}(\gamma) \leq k$ defines D, then the set $D = \xi^{-1}[\{\gamma^{\mathfrak{M}} > 0\}]$ of the proof of Theorem 7.7 is $\mathbf{\Delta}^1_{k+1}$ by part (1). For part (3), the set C is not projective, so by part (2) every formula that defines it has quantifier depth at least ω. As the subformula relation is well founded, such a formula has a subformula of depth at least ω whose proper subformulas all have finite depth. This subformula is a countable join whose members have unbounded finite depths. □

Uniformity across standard models. Definability over T requires one formula to work independently of the basis, the density operator and the unit representatives chosen by rep (see Definition 7.6). In the model $\mathfrak{M}$ of the proof of Theorem 7.7, whose symbols are interpreted by Borel functions, definability by one formula in this single model already gives definability over T. Indeed, let $D \subseteq [0,1]$ be a set such that some formula takes, at every element Q of the Solèr sort of $\mathfrak{M}$, the value 1 if $\mu(Q) \in D$ and -1 otherwise. The second half of the proof applies to this formula and shows that the set D is σ-projective, and hence definable over T.

In an arbitrary standard model, definability by one formula cannot replace definability over T. For $0 < t < 1$, the ray of the unit vector $\sqrt{t}\, e_0 + \sqrt{1-t}\, e_1$ has probability t in the pure state of e_0, and the choice of rep between this vector and its negative can encode an arbitrary subset of $(0,1)$ in one standard model. We detect the choice by requiring both coordinates of the chosen representative to be positive and real.

To express this test, for $i \in \mathbb{N}$ we write $z_i(b)$ for the pair of predicates $\langle \mathrm{rep}(b) \mid e_i \rangle$, and $\mathrm{nonatom}(b)$ for the two-valued formula $\sup_c\big([c \neq 0] \wedge -[c = b] \wedge [c \leq b]\big)$, which takes the

value 1 exactly when some element lies strictly between 0 and b. For an atomic formula $R(\bar{t})$ whose predicate symbol is one of the inner-product, norm or measure predicates, let $[R(\bar{t}) > 0] = 2\lceil R(\bar{t}) > 0\rceil - 1$ and $[R(\bar{t}) = 0] = 1 - 2\big(\lceil R(\bar{t}) > 0\rceil \vee \lceil R(\bar{t}) < 0\rceil\big)$, two-valued formulas in every structure that satisfies (Tr). Let $\mathrm{eq}_\mu(a,b)$ be the formula $\beta_B(a,b)$ of the proof of Theorem 7.7 for the diagonal $B = \{(u,u) : u \in \mathbb{R}\}$. Let $\gamma_0(a)$ be the formula

$$\sup_b \Big([b \neq 0] \wedge -\mathrm{nonatom}(b) \wedge \big[b \leq \mathrm{at}(e_0) \vee \mathrm{at}(e_1)\big] \wedge \mathrm{eq}_\mu(a,b) \\ \wedge [\operatorname{Im} z_0(b) = 0] \wedge [\operatorname{Im} z_1(b) = 0] \wedge [\operatorname{Re} z_0(b) > 0] \wedge [\operatorname{Re} z_1(b) > 0] \Big).$$

Proposition 7.9 (Definability in one standard model is not uniform). *For every set $\Omega \subseteq (0,1)$, there is a standard model of T in which the formula $\gamma_0(a)$ takes the value 1 exactly when $\mu(a) \in \Omega$, and -1 otherwise. Consequently, some sets that the single formula γ_0 defines in a single standard model are not definable over T.*

Proof. Given $\Omega \subseteq (0,1)$, take the model of the proof of Theorem B(2) on $\ell^2(\mathbb{N})$, with the standard unit vectors as the constants e_n and with ρ the projection onto $\mathbb{C}e_0$. In this model, let rep send the atom spanned by $\sqrt{t}\,e_0 + \sqrt{1-t}\,e_1$, for $0 < t < 1$, to the vector $\sqrt{t}\,e_0 + \sqrt{1-t}\,e_1$ itself if $t \in \Omega$ and to its negative otherwise. For $0 < t < 1$, the atoms b below $\mathrm{at}(e_0) \vee \mathrm{at}(e_1)$ with $\mu(b) = t$ are the lines of the vectors $\sqrt{t}\,e_0 + \zeta\sqrt{1-t}\,e_1$ with $|\zeta| = 1$. A unit vector of such a line has two real coordinates only when $\zeta = \pm 1$, and two positive real coordinates only when $\zeta = 1$ and the vector is $\sqrt{t}\,e_0 + \sqrt{1-t}\,e_1$ itself; for $t \in \{0,1\}$, one coordinate vanishes. Hence the formula $\gamma_0(a)$ takes the value 1 in this model exactly when $\mu(a) \in \Omega$, and -1 otherwise.

Since the vocabulary is countable, there are at most $2^{\aleph_0}$ formulas, and hence at most $2^{\aleph_0}$ sets definable over T. The interval $(0,1)$, by contrast, has $2^{2^{\aleph_0}}$ subsets, and some sets Ω are therefore not definable over T. □

A quantified formula defining a set D over T need not take the value 1 exactly when $\mu(a) \in D$ in a nonstandard model. By Remark 10.3, the theory T has a countable model $\mathfrak{N}$, so some $\alpha \in [0,1]$ is absent from its countable range of μ. The formula $\sup_b \beta_{\mathbb{R}\times\{\alpha\}}(a,b)$, with the Borel test constructed in the proof of Theorem 7.7, searches for an element of measure α. It defines $[0,1]$ over T by Proposition 7.5, yet takes the value -1 everywhere in $\mathfrak{N}$.

The quantifier-free tests $\beta_B(a)$ of Lemma 4.1 retain their intended meaning in every structure satisfying (Tr): they take the value 1 exactly when $\mu(a) \in B$. Remark 8.8 below concerns all models of T and therefore uses these Borel tests. The classification in Theorem 7.7 concerns standard models, where the full interval of probabilities is available to the quantifiers.

8. The proof of Theorem C

The preceding representation theorems concern standard models, but the exclusion of a $\{0,1\}$-valued measure holds in every model of T. We prove Theorem C directly from the axioms, using the Hilbert space $\bar{E}$ constructed in Lemma 5.3 to supply the geometry. Assuming that μ takes only the values 0 and 1, we apply (M5) to the named sequence ε to locate a basis atom of measure 1, then obtain a finite parity contradiction within a block of four basis vectors.

Within the model, the argument uses only the vectors denoted by the closed terms described below. By (H2), the constants e_n are orthonormal, and by (K4) the entries of the sequence constant ε satisfy $\kappa_n \varepsilon \doteq 2^{-n-1} e_n$. These equations use the relation $\doteq$ of Section 4, which asserts that a scaled difference has norm zero. By Lemma 5.3, the related elements

have the same image in $\bar{E}$. They need not be identical as elements of the model, since the equality-free vocabulary imposes no requirement that elements agreeing in all predicates be identical. Closed terms formed from four of the constants e_n using the functions $\lambda x + \lambda' y$ represent nonzero scalar multiples of the eighteen vectors of Table 1, and therefore their rays. Their images in $\bar{E}$ suffice for the parity argument, which requires neither the whole unit ball nor the whole lattice of closed subspaces inside the model.

The localization step needs the named orthonormal basis and the measure of its span, while the parity step needs the corresponding rays and additivity on orthogonal subspaces. The following lemma gathers these facts in a form valid for every model of T. For a model of T_1, we let $\bar{E}$, the map $x \mapsto \hat{x}$ and the projections $\bar{P}_a$ be as in Lemma 5.3, applied to the reduct of the model to the vocabulary of T_0, and we let S_a be the range of $\bar{P}_a$, as in Lemma 5.4.

Lemma 8.1. *Let $\mathfrak{M}$ be a model of T, standard or not, and let $\bar{E}$, $\hat{x}$ and S_a be as above. Then the following hold:*

(1) The vectors $\hat{e}_n$, for $n \in \mathbb{N}$, form an orthonormal basis of $\bar{E}$.
(2) For every x in the Hilbert sort with $\|x\| > 0$, the subspace $S_{\mathrm{at}(x)}$ is $\mathbb{C}\hat{x}$, and the formula $[\mathrm{at}(x) = \mathrm{at}(y)]$ takes the value 1 whenever $\|x\| > 0$ and $\hat{y} = \lambda\hat{x}$ for a nonzero $\lambda \in \mathbb{C}$.
(3) The measure satisfies $\mu(0) = 0$, and $\mu(a \vee c) = \mu(a) + \mu(c)$ for all elements a, c of the Solèr sort with $S_c \subseteq S_a^{\perp}$.
(4) We have $[\mathrm{span}(\varepsilon) = 1] = 1$ and $\mu(\mathrm{span}(\varepsilon)) = 1$.

Proof. (1) By (H2), the vectors $\hat{e}_n$ are orthonormal. By Lemma 5.4, the projection $\bar{P}_{b_N}$ is the projection onto the span of $\hat{e}_0, \dots, \hat{e}_{N-1}$, and (H2) places every vector $\hat{x}$ in the closed span of the vectors $\hat{e}_n$; since the vectors $\hat{x}$ span a dense subspace, the vectors $\hat{e}_n$ form an orthonormal basis of $\bar{E}$.

(2) The first assertion is part (4) of Lemma 5.4. If $\hat{y} = \lambda\hat{x}$ with $\lambda \neq 0$, then $\|y\| > 0$ and $S_{\mathrm{at}(y)} = \mathbb{C}\hat{x} = S_{\mathrm{at}(x)}$, and part (3) of the same lemma gives $[\mathrm{at}(x) = \mathrm{at}(y)] = 1$.

(3) The argument of Step 1 of the proof of Theorem B applies to $\mathfrak{M}$, with S_a the range of $\bar{P}_a$, because it uses, besides axioms (M3) and (M4), only Lemma 5.4, which holds in every model of T_0.

(4) By (Sp1) and (K4), the subspace $S_{\mathrm{span}(\varepsilon)}$ contains every vector $2^{-n-1}\hat{e}_n$, and hence, by part (1), it is $\bar{E} = S_1$. Part (3) of Lemma 5.4 gives $[\mathrm{span}(\varepsilon) = 1] = 1$, and axioms (M3) and (M2) give $\mu(\mathrm{span}(\varepsilon)) = 1$. □

The proofs of parts (1) and (2), and of the equation $[\mathrm{span}(\varepsilon) = 1] = 1$ in part (4), use no axiom that mentions μ; these assertions therefore hold in every model of T_1.

Proof of Theorem C. Suppose, towards a contradiction, that $\mathfrak{M}$ is a model of T in which μ takes only the values 0 and 1, and let $\bar{E}$, $\hat{x}$ and S_a be as in Lemma 8.1 for this model. We first show that exactly one of the constants e_n has an atom $\mathrm{at}(e_n)$ of measure 1. We then apply a parity argument to the bases of Table 1 within the block of four consecutive constants that contains this one. The entries $\kappa_n\varepsilon$ satisfy $\widehat{\kappa_n\varepsilon} = 2^{-n-1}\hat{e}_n$ by (K4), and part (2) of Lemma 8.1 gives $[\mathrm{at}(\kappa_n\varepsilon) = \mathrm{at}(e_n)] = 1$. Since the subspaces $S_{\mathrm{at}(e_n)} = \mathbb{C}\hat{e}_n$ are mutually orthogonal, repeated applications of part (3) of Lemma 8.1, combined with (M3), which replaces each $\mu(\mathrm{at}(\kappa_n\varepsilon))$ by $\mu(\mathrm{at}(e_n))$, give $\mu(c_N(\varepsilon)) = \sum_{n<N} \mu(\mathrm{at}(e_n))$ for every $N \in \mathbb{N}$. Since the terms of these partial sums are nonnegative by (M1), axiom (M5) and Lemma 8.1(4) give

$$1 = \mu(\mathrm{span}(\varepsilon)) = \sum_n \mu(\mathrm{at}(e_n)).$$

basis	vectors (coordinates with respect to an orthonormal basis g_0, g_1, g_2, g_3)			
1	$(0,0,0,1)$	$(0,0,1,0)$	$(1,1,0,0)$	$(1,-1,0,0)$
2	$(0,0,0,1)$	$(0,1,0,0)$	$(1,0,1,0)$	$(1,0,-1,0)$
3	$(1,-1,1,-1)$	$(1,-1,-1,1)$	$(1,1,0,0)$	$(0,0,1,1)$
4	$(1,-1,1,-1)$	$(1,1,1,1)$	$(1,0,-1,0)$	$(0,1,0,-1)$
5	$(0,0,1,0)$	$(0,1,0,0)$	$(1,0,0,1)$	$(1,0,0,-1)$
6	$(1,-1,-1,1)$	$(1,1,1,1)$	$(1,0,0,-1)$	$(0,1,-1,0)$
7	$(1,1,-1,1)$	$(1,1,1,-1)$	$(1,-1,0,0)$	$(0,0,1,1)$
8	$(1,1,-1,1)$	$(-1,1,1,1)$	$(1,0,1,0)$	$(0,1,0,-1)$
9	$(1,1,1,-1)$	$(-1,1,1,1)$	$(1,0,0,1)$	$(0,1,-1,0)$

TABLE 1. The configuration of Cabello, Estebaranz and García-Alcaine [CEGA96]: nine orthogonal bases of a four-dimensional space, formed by eighteen vectors, each of which occurs in exactly two bases.

As μ takes only the values 0 and 1, there is exactly one $n_0 \in \mathbb{N}$ with $\mu(\mathrm{at}(e_{n_0})) = 1$. Let $j \in \mathbb{N}$ satisfy $4j \leq n_0 < 4j+4$, write $g_i = e_{4j+i}$ for $i = 0,1,2,3$, and let F be the span of $\hat{g}_0, \ldots, \hat{g}_3$.

Among the four constants $g_0, \ldots, g_3$, whose images $\hat{g}_0, \ldots, \hat{g}_3$ form an orthonormal basis of F, the constant e_{n_0} is the only one whose atom has measure 1. We show that every orthogonal basis of F whose vectors have coordinates in $\{0, 1, -1\}$ with respect to $\hat{g}_0, \ldots, \hat{g}_3$ likewise contains exactly one vector v whose atom $\mathrm{at}(v)$, defined below, has measure 1. Let $\mathcal{B}$ be such a basis, consisting of vectors $v_0, \ldots, v_3$ of F. The coordinates of each vector $\frac{1}{4}v_i$ have absolute values with sum at most 1; by Lemma 5.3, some closed term t_i, built from the constants $g_0, \ldots, g_3$ by the functions $\lambda x + \lambda' y$, therefore satisfies $\hat{t}_i = \frac{1}{4}v_i$. Hence the term $\frac{1}{2}\big(\frac{1}{2}\iota_0 t_0 + \frac{1}{2}\iota_1 t_1\big) + \frac{1}{2}\big(\frac{1}{2}\iota_2 t_2 + \frac{1}{2}\iota_3 t_3\big)$ denotes an element $s_{\mathcal{B}}$ of the sequence sort. Axiom (K2) gives $\langle \kappa_m \iota_n x \mid y\rangle = \langle \iota_n x \mid \iota_m y\rangle = \delta_{mn}\langle x \mid y\rangle$ for all x and y, whence $\widehat{\kappa_m \iota_n x} = \delta_{mn}\hat{x}$, because the vectors $\hat{y}$ span a dense subspace of $\bar{E}$; together with the linearity of κ_n asserted in (K2), this gives $\widehat{\kappa_n s_{\mathcal{B}}} = \frac{1}{16}v_n$ for $n < 4$. Let $c_{\mathcal{B}}$ be the element $\mathrm{at}(\kappa_0 s_{\mathcal{B}}) \vee \cdots \vee \mathrm{at}(\kappa_3 s_{\mathcal{B}})$, that is, $c_4(s_{\mathcal{B}})$.

For a nonzero vector v of F with coordinates in $\{0, 1, -1\}$, we fix a closed term t_v, built in the same way, with $\hat{t}_v = \frac{1}{16}v$, and we write $\mathrm{at}(v)$ for $\mathrm{at}(t_v)$. By Lemma 8.1(2), another choice of t_v changes $\mathrm{at}(v)$ only within the relation $[\cdot = \cdot] = 1$, and therefore, by (M3), leaves $\mu(\mathrm{at}(v))$ unchanged.

By the same part of Lemma 8.1, we have $[\mathrm{at}(\kappa_i s_{\mathcal{B}}) = \mathrm{at}(v_i)] = 1$ for $i < 4$. By Lemma 5.4, the subspace $S_{c_{\mathcal{B}}}$ is the span of $v_0, \ldots, v_3$, which is F. Hence $[c_{\mathcal{B}} = c_{\mathcal{B}'}] = 1$ for any two such bases $\mathcal{B}$ and $\mathcal{B}'$, and we let a_F be the element $c_{\mathcal{B}}$ for the basis $\hat{g}_0, \ldots, \hat{g}_3$. By Lemma 8.1(3), applied three times, and by (M3),

$$\mu(a_F) = \mu(c_{\mathcal{B}}) = \sum_{v \in \mathcal{B}} \mu(\mathrm{at}(v)).$$

For the basis $\hat{g}_0, \ldots, \hat{g}_3$, part (2) of Lemma 8.1 and (M3) give $\mu(\mathrm{at}(\hat{g}_i)) = \mu(\mathrm{at}(e_{4j+i}))$ for $i < 4$, and this sum therefore contains $\mu(\mathrm{at}(e_{n_0})) = 1$ and three zeros; hence $\mu(a_F) = 1$, and every such basis $\mathcal{B}$ contains exactly one vector v with $\mu(\mathrm{at}(v)) = 1$.

The nine bases of Table 1 are orthogonal bases of F of this kind, and their eighteen vectors span pairwise distinct rays. We conclude by counting in two ways the pairs $(\mathcal{B}, v)$ that consist of one of the nine bases $\mathcal{B}$ and a vector v of $\mathcal{B}$ with $\mu(\mathrm{at}(v)) = 1$. By the preceding paragraph, each basis contains exactly one such vector, and there are therefore exactly nine pairs. On the other hand, each of the eighteen vectors v occurs in exactly two of the bases, and the atom $\mathrm{at}(v)$ fixed above is the same element in both. Each vector v with $\mu(\mathrm{at}(v)) = 1$ thus belongs to exactly two pairs, and the number of pairs is even. This contradiction proves the theorem. □

In physical terms, Theorem C excludes noncontextual hidden-variable assignments with probabilities in $\{0, 1\}$ that satisfy the measure axioms of T. The experimentally verifiable propositions are represented by closed subspaces [BvN36], whose role in a model of T is played by the elements of the Solèr sort. For each value of the hidden variable, the assignment gives a definite truth value to each proposition, independently of which other propositions are tested together with it. Such an assignment would be recorded by a dispersion-free measure, taking the numerical values 0 and 1 for false and true respectively. By Theorem C, the axioms of T, including the continuity requirement (M5) along represented spans, exclude this assignment as the interpretation of μ. Contextual theories, such as Bohm's [Boh52], allow the value revealed by a measurement to depend on the whole experimental arrangement and lie outside this conclusion. Together with (M1), the theorem gives the following corollary.

Corollary 8.2 (No dispersion-free measures). *In every model of T, standard or not, some element a of the Solèr sort satisfies $0 < \mu(a) < 1$.*

Proof. By (M1), the measure takes values in $[0, 1]$, and by Theorem C it does not take only the values 0 and 1. □

⧖ **Historical note.** Specker announced in 1960 that the closed subspaces of a three-dimensional Hilbert space cannot be mapped into a Boolean algebra in a way that respects the operations on orthogonal subspaces. Of the proof he said only that it rests on an elementary geometric argument [Spe60]. Bell derived from Gleason's theorem that no hidden-variable theory of the kind described above is possible, in a paper written in 1964 and published in 1966 [Bel66].[5] In 1967 Kochen and Specker gave a finite proof of this impossibility with 117 vectors in $\mathbb{R}^3$ [KS67]. The configurations then shrank: Peres found configurations of 33 vectors in $\mathbb{R}^3$ and of 24 in $\mathbb{R}^4$ in 1991 [Per91], and Kernaghan one of 20 vectors in $\mathbb{R}^4$ in 1994 [Ker94]. In 1996 Cabello, Estebaranz and García-Alcaine found the eighteen vectors of Table 1 [CEGA96]. Peres conjectured that no smaller configuration exists, and Xu, Chen and Gühne proved in 2020 that no Kochen–Specker set in any dimension has fewer than eighteen vectors [XCG20, BCG+22]. In dimension three the minimum remains unknown at the time of writing. Exhaustive computer searches raised the lower bound for this minimum to 22 vectors in 2016 [UW16] and to 24 in 2023. The bound of 24 was reached independently by Kirchweger, Peitl and Szeider and by Li, Bright and Ganesh [KPS23, LBG24]. The titles of the papers of Kernaghan and of Cabello, Estebaranz and García-Alcaine call the result the Bell–Kochen–Specker theorem, a name that records Bell's share and that Theorem C keeps.

[5] Bell's geometric comparisons of vectors and rays [Bel66, Section V] admit a modern topological interpretation. For unit vectors f_0, f_1 of a complex Hilbert space, let Q_i be the orthogonal projection onto $\mathbb{C}f_i$, for $i = 0, 1$. The identities $\|Q_0 - Q_1\| = \|(Q_0 - Q_1)f_0\| = \sqrt{1 - |\langle f_0 \mid f_1\rangle|^2}$ and $\inf_{|\lambda|=1} \|f_0 - \lambda f_1\| = \sqrt{2 - 2|\langle f_0 \mid f_1\rangle|}$, with $\lambda \in \mathbb{C}$, identify both the strong operator topology and the operator norm topology on rank-one projections with the quotient of the norm topology on the unit sphere by scalar phases.

8.1. **The Bell–Kochen–Specker obstruction in the game.** The contradiction of Theorem C gives a concrete case in which the finite obstructions of Section 3 arise from an infinitary theory. In this subsection the vocabulary and the theory are again those of Section 4 and of Theorem B. We extend the theory T of Theorem B by the assertion that its measure μ takes only the values 0 and 1. The normal form of this assertion is $\sup_a\big(\mu(a) \wedge (\mathbf{1} - \mu(a))\big) \leq 0$, which we call the *dispersion-free axiom*. By Theorem C, the extended theory has no models, so Corollary 3.5 applies. In Proposition 8.3 below we show that every finite set of its axioms has a model, and in Corollary 8.4 we deduce that the tree of every winning strategy of Olivia in the game for this theory has an infinitely branching node. In Proposition 8.6 below we construct a strategy following the proof of Theorem C: every play reaching the parity argument ends with commitments concerning the eighteen vectors and nine bases of Table 1 in a block of four consecutive constants e_n.

Proposition 8.3. *Let T^{01} be the theory T of Theorem B together with the axiom*

$$\sup_a\big(\mu(a) \wedge (\mathbf{1} - \mu(a))\big) \leq 0.$$

(1) Olivia has a winning strategy in $\mathrm{G}(T^{01})$; every play in which she follows it ends after finitely many rounds, at a position that is satisfied in no structure.
(2) Every finite set of axioms of T^{01} has a model.
(3) The skeleton of every winning strategy of Olivia in $\mathrm{G}(T^{01})$ has more than one element: in some play in which Olivia follows such a strategy, Josh answers a countable challenge.

Proof. Part (1) follows from parts (1) and (2) of Corollary 3.5, since T^{01} has no models by Theorem C.

For part (2), recall that in the proof of Theorem C it is the instances of (K4), taken together, that make the entries of ε span the whole space, and that only then does the σ-additivity of μ along ε give some atom the measure 1. A finite set of axioms contains only finitely many of these instances, and we therefore look for a model of it among the structures whose sequences are all finitely supported. In such a structure, axiom (M5) involves only finite-dimensional subspaces, and a measure that vanishes on every finite-dimensional subspace satisfies it. Let a finite set of axioms of T^{01} be given, and let $Z \subseteq \mathbb{N}$ be the finite set of the numbers n for which this set contains the instance $\kappa_n\varepsilon \doteq 2^{-n-1}e_n$ of axiom (K4). A model of the given set is the following structure. Its Hilbert sort is the closed unit ball of the space ℓ^2 of square-summable complex sequences, and its Solèr sort is the lattice L_{fc} of Example 6.5. The symbols π, at, rep and the constants e_n are interpreted as in this example. The sequence sort consists of the finitely supported sequences in the closed unit ball of the space $\ell^2(\mathbb{N}, \ell^2)$ of square-summable sequences in ℓ^2. The interpretations of ι_n, κ_n and res_N are those of the proof of Theorem 7.3, and the function span sends each element s of the sequence sort to the finite-dimensional span of its entries. The measure μ vanishes on the finite-dimensional subspaces and equals 1 on their orthogonal complements, and each truncation $R^{\pm}_{n,q}$ is interpreted by its intended value $\min\big(1, \max(0, \pm n(R-q))\big)$, so that the axioms (Tr) hold. The constant ε is the finitely supported sequence whose entry of index n is $2^{-n-1}e_n$ for $n \in Z$ and 0 otherwise; its norm is less than 1. The axioms of T_0 hold as in Example 6.5, and the axioms (K1)–(K3) and (Sp) hold as in the proof of Theorem 7.3, since the span of finitely many vectors is closed; in this structure, the formula $[a = b]$ takes the value 1 exactly when $a = b$. Of the instances of (K4), the only axiom in which ε occurs, exactly those with $n \in Z$ hold. Axiom (M3) follows from the description of the formula $[a = b]$ given above, and axioms (M1), (M2) and the dispersion-free axiom hold because μ takes only the values 0 and 1, with $\mu(1) = 1$. For (M4),

the subspaces a and $b \wedge a^{\perp}$ are orthogonal, and at most one of them is infinite-dimensional; the measure of their join is 1 exactly when one of them has measure 1. For (M5), let s be an element of the sequence sort; the element $\mathrm{span}(s)$ and the elements $c_N(s)$ are finite-dimensional subspaces, and both sides of (M5) vanish. The structure is therefore a model of the given finite set of axioms.

For part (3), suppose that the skeleton of a winning strategy σ of Olivia had only one element, the empty sequence. Then Josh would give no countable answer at any node of $\mathcal{P}_\sigma$: a node at which σ makes a countable challenge would have only one child, the extension by the challenge alone. Every other node would have finitely many children by Corollary 3.5(3). By Corollary 3.5(4), the theory T^{01} would then have a finite set of axioms without models, contradicting part (2). □

Corollary 8.4. *Every winning strategy of Olivia in the model existence game for the theory of Theorem C has a node at which its tree branches infinitely.*

Proof. The theory of Theorem C, written in the normal form of Subsection 3.1, is the theory T^{01} of Proposition 8.3. By part (2) of the proposition, every finite set of axioms of T^{01} has a model. It follows from part (3) that Josh gives a countable answer at some node of the tree of every winning strategy of Olivia. The interval that Josh adds under the rules $(\bigvee \geq)$ and $(\bigwedge \leq)$ does not depend on the number he names. Every number is therefore a legal answer at this node, and the tree branches infinitely there. □

The proof of Theorem C first uses countable additivity to locate an atom of measure 1, then derives a finite parity contradiction in the block containing it. To follow the proof in the game, Olivia needs recorded bounds with enough numerical margin for the next step. For example, the equation $\mu(\mathrm{span}(\varepsilon)) = 1$ uses every instance of (K4), whereas any position records only finitely much information. We therefore first explain how Olivia can observe values to arbitrary accuracy. In Lemma 8.5 below we show how fair bookkeeping either exposes a contradiction or obtains the required margins after finitely many rounds in each play, without asserting a uniform bound on their number.

Observing values. The margins in the localization and parity argument concern predicate values at closed terms. Olivia first names the required closed terms, from their innermost subterms outward, by finitely many challenges (Tot); the resulting records make the names denote the terms in every structure satisfying the position. She then obtains intervals of arbitrarily small width for their predicate values, using the truncation axioms. For a predicate symbol R other than $\mathbf{1}$ and the truncation symbols, the procedure at given witnesses is as follows. Challenges (Ax) activate the axiom (Tr) for R with $n = 1$ and $q = 0$; challenges $(\sup \leq)$ and $(\inf \geq)$ then give commitments on its instance at the given witnesses. Olivia follows the finite chain of connective functions down to the occurrence of R by challenges (Cn), ending with the addition $R(\bar{w}) - 0$. The interval $I_{R(\bar{w})}$ has length at least 2, since the a priori bound of the predicate symbol is a positive integer, whereas the interval of the sentence -0 is the single point 0. The number m prescribed by (Cn) for this addition is therefore k, so a tolerance 2^{-k} yields an interval of width at most 2^{-k} for $R(\bar{w})$. These challenges may expose a contradiction before the observation is complete, in which case Josh loses. By an observation of $R(\bar{t})$ we mean the naming records together with an interval commitment on the resulting basic sentence $R(\bar{w})$; bounds on the value of $R(\bar{t})$ follow from the records.

At the start of the play, Olivia makes the challenge (Tot) for the constant ε, and we write w_ε for the witness of the record $\varepsilon \mapsto w_\varepsilon$ that Josh adds. Below, Olivia also observes the value

of the sentence

$$\chi = \bigvee_{N\in\mathbb{N}} \mu(c_N(w_\varepsilon)).$$

The instance of (M5) at w_ε, obtained by (Ax) followed by the rules ($\sup\leq$) and ($\inf\geq$), contains χ as an argument of the negation in $\mu(\mathrm{span}(w_\varepsilon)) - \chi$. Challenges (Cn) along this finite chain give intervals of arbitrarily small width for χ as well. Here the observation concerns the value of the join, not the choice of one of its members; such a choice will be made separately by ($\bigvee\geq$).

Lemma 8.5 (Finite continuations). *Let S be a nonempty countable subtheory of T^{01}, regarded as a theory in the full vocabulary, and let $\mathfrak{p}$ be a position reached in the game. Let $\psi_1, \dots, \psi_M$, for a positive integer M, be sentences whose values can be observed as above using axioms of S. Let $U \subseteq \mathbb{R}^M$ be open, and suppose that every model of S satisfying $\mathfrak{p}$, under an interpretation of its witnesses, has*

$$(\psi_1, \dots, \psi_M) \in U.$$

Olivia has a continuation, using only axioms of S, under which every play either ends with Josh losing or reaches, after finitely many rounds, a position whose commitments confine this tuple of values to a nonempty closed box contained in U.

Proof. Fix an enumeration of all the challenges that can occur in continuations of $\mathfrak{p}$ in which every challenge (Ax) names an axiom of S; among them are the challenges (Tot), the challenges of the observing procedures and the challenges to the commitments of $\mathfrak{p}$. At stage $k \geq 1$, Olivia makes every challenge among the first k entries that is available at the start of the stage, and then observes all the listed values at tolerance 2^{-k}, in a fixed order. Each stage has finitely many rounds unless Josh loses, and every challenge that becomes available is eventually made, since availability persists. For each observed sentence she intersects its recorded intervals with its fixed interval of possible values; these intersections form nested closed intervals, and the observations make their widths tend to zero. She ends the continuation when their product is nonempty and contained in U.

Suppose that a play continued indefinitely without either outcome. It would be fair for S, and the construction in the proof of Theorem 3.3, applied from the initial position $\mathfrak{p}$, would give a model of S satisfying every commitment and record made in the play. This construction does not require the initial position to be empty: its commitments are included in the same induction on formulas, and its records in the definition of the functions. The tuple of observed values in this model belongs to U by hypothesis and lies in every recorded box. Since the boxes have widths tending to zero and U is open, some box is contained in U, a contradiction. An empty intersection cannot occur in such an infinite surviving play, since the constructed model supplies a point of every interval. Thus every play has one of the asserted finite outcomes. □

The openness of the target lets Olivia stop when the recorded intervals themselves give the required margin, for example when a value exceeds a rational threshold or a finite sum lies within a prescribed tolerance of another value. For the targets chosen in localization and parity, the consequences of Lemmas 5.3, 5.4 and 8.1 verify the hypothesis of Lemma 8.5: every model of the chosen subtheory S satisfying the current position has its tuple of values in the target open set. The lemma converts these conclusions about all such models into finite commitments with the required margins. At no stage must Olivia decide whether the current position has a model of S: she stops when its recorded box lies in U, or wins if Josh loses first (see Figure 7).

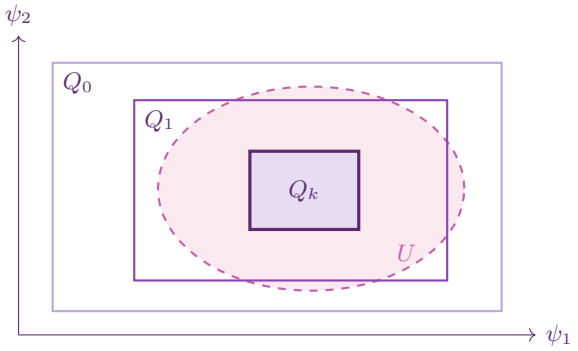


FIGURE 7. The finite continuation of Lemma 8.5, drawn for two observed values ψ_1 and ψ_2, in which Olivia stops at the first nonempty box Q_k of recorded intervals contained in the open set U.

Proposition 8.6 (The localization and parity strategy). *Olivia has a winning strategy in* $\mathrm{G}(T^{01})$ *such that every play in which she follows it ends after finitely many rounds and, unless Josh loses earlier, passes through the following two phases:*

(1) Josh gives a countable answer $N \geq 1$*, and Olivia selects an index* $n_0 < N$ *for which Josh has committed the measure of the atom* $\mathrm{at}(\kappa_{n_0}\varepsilon)$ *to an interval whose left endpoint exceeds* $15/16$*;*

(2) the play records finite geometric commitments for the nine bases of Table 1 in the block containing e_{n_0}*, obtained without any axiom that mentions* μ*, and concludes by a finitely branching continuation that uses finitely many instances of (M1), (M3), (M4) and the dispersion-free axiom.*

Proof. Localization. In every model of T satisfying the record for w_ε, the sentence χ has value 1, by (M5) and Lemma 8.1(4). Apply Lemma 8.5 with $S = T$ and the open interval $(15/16, 17/16)$. Unless Josh loses first, the intersection of the recorded intervals for χ has lower endpoint greater than $15/16$. This endpoint is the maximum of finitely many recorded lower endpoints and the fixed lower bound -2; hence one actual commitment for χ has lower endpoint $\alpha > 15/16$. A challenge $(\bigvee \geq)$ with tolerance $1/16$ then makes him name an index $N \in \mathbb{N}$ and commit

$$\mu(c_N(w_\varepsilon)) \geq \alpha - 1/16 > 7/8.$$

If $N = 0$, the selected term is $c_0(w_\varepsilon) = 0$; fair bookkeeping for T exposes a contradiction after finitely many rounds, since an infinite surviving play would give a model of T with $\mu(0) > 7/8$, contrary to Lemma 8.1(3). We may therefore suppose that $N \geq 1$.

Write $h_n = \mathrm{at}(\kappa_n w_\varepsilon)$ for $n < N$. In every model of T satisfying the record $\varepsilon \mapsto w_\varepsilon$, the finite-additivity calculation in the proof of Theorem C gives

$$\sum_{n<N} \mu(h_n) = \mu(c_N(w_\varepsilon)).$$

In every model of T satisfying the current position, the commitment just made gives $\mu(c_N(w_\varepsilon)) > 7/8$, and the sum of the values $\mu(h_n)$ therefore exceeds $3/4$. A further application of Lemma 8.5, with $S = T$ and the open set of tuples whose sum exceeds $3/4$, confines the values $\mu(h_n)$ to a closed box on which their sum exceeds $3/4$. Let $\delta_N = \min(1/16, 1/(4N))$. We apply Lemma 8.5 once more, now with S consisting of (M1), the dispersion-free axiom and the axioms (Tr) for μ, and with the open set of tuples whose coordinates all lie in

$(-\delta_N, \delta_N) \cup (1-\delta_N, 1+\delta_N)$. In every model of these axioms, each value $\mu(h_n)$ is 0 or 1, and the hypothesis of the lemma is therefore satisfied. The lemma gives a refinement of the recorded box in which each coordinate interval lies within one of these two intervals. If all intervals were of the first kind, their sum would be less than $N\delta_N \le 1/4$, contrary to the preceding box condition. Choose the least index $n_0 < N$ whose interval is of the second kind. The lower endpoint of the interval for $\mu(h_{n_0})$ exceeds $1-\delta_N \ge 15/16$; as for χ, it is the lower endpoint of one of the commitments recorded for $\mu(h_{n_0})$, since it exceeds the fixed lower bound -2. This completes the first phase.

Geometric identities. Choose $j \in \mathbb{N}$ with $4j \le n_0 < 4j+4$ and use the closed terms $s_{\mathcal{B}}$ and $\mathrm{at}(v)$ of the proof of Theorem C for the nine bases of Table 1 in the block containing e_{n_0}. The four contextual atom terms of a basis $\mathcal{B}$ are $a_{\mathcal{B},i} = \mathrm{at}(\kappa_i s_{\mathcal{B}})$, for $i < 4$. We consider, in addition, the basis of the block formed by the four constants $e_{4j}, \dots, e_{4j+3}$, ordered with e_{n_0} first and the other three in increasing order of index. For this basis we use the atom terms $\mathrm{at}(e_i)$ themselves, rather than contextual ones, and we now let a_F be their join in this order; like the element a_F of the proof of Theorem C, it names the span of the block. To apply (M4) successively, we replace each new atom by its meet with the orthocomplement of the preceding partial join. For any of these ten ordered bases, write its atom terms as $a_0, \dots, a_3$, and form the terms

$$d_1 = a_0, \qquad d_{i+1} = d_i \vee a_i, \qquad a'_i = a_i \wedge d_i^{\perp}, \qquad d'_{i+1} = d_i \vee a'_i \quad (1 \le i \le 3).$$

Let $\mathcal{D}$ be the following finite list of pairs of closed terms of the Solèr sort: the two contextual terms of each shared ray; the pairs (a'_i, a_i) and (d'_{i+1}, d_{i+1}) for each ordered basis; the pairs (d_4, a_F) for these bases; and the pair $(h_{n_0}, \mathrm{at}(e_{n_0}))$. Olivia will obtain commitments making the formula $[u = v]$ positive for each pair $(u, v) \in \mathcal{D}$, so that (M3) forces the two measure values to agree. In every model of T_1 satisfying the records, the two terms of each pair (u, v) in this list name the same closed subspace, and they therefore satisfy $[u = v] = 1$ by part (3) of Lemma 5.4. For the atom terms, these subspaces are computed as in the proof of Theorem C, by Lemma 5.3, (K2), (K4) and part (2) of Lemma 8.1; for the other terms, they follow from those of the atom terms by part (1) of Lemma 5.4. These conclusions use no measure axiom.

For such a pair, the formula $[u = v]$ is the minimum of $1 - 2\|\mathrm{rep}(u \wedge (u \wedge v)^{\perp})\|$ and the corresponding expression with u and v interchanged. The norms in these expressions vanish in every model of T_1 satisfying the records. Apply Lemma 8.5 with $S = T_1$ to these finitely many atomic norms and the open set of tuples in which each of these norms is less than $1/8$. Unless Josh loses earlier, the resulting basic commitments ensure

$$[u = v] > 3/4 \qquad ((u, v) \in \mathcal{D})$$

in every structure satisfying the position. Name the finitely many closed subterms involved by (Tot), retaining their records. Thus both the contextual ray identities and the identities needed for finite additivity have been recorded to a definite margin before any probability is carried across them.

The finite parity continuation. Retain the finite set $\mathfrak{p}_0$ consisting of these function records and norm commitments, together with the interval for $\mu(h_{n_0})$ whose lower endpoint exceeds $15/16$. Bring into play the instances of (M3) at the pairs in $\mathcal{D}$, the instances of (M4) at (d_i, a_i) for the ten bases, and the instances of (M1) and the dispersion-free axiom at the finitely many terms occurring in these measure conditions. Each instance is obtained by (Ax) followed by the rules (sup $\le$) and (inf $\ge$), applied to the witnesses that name its terms; an equation contributes its two closed inequalities. Denote the resulting finite family of ground

commitments by $\mathfrak{q}$. All its sentences are built from atomic sentences by the connective functions alone, without suprema, infima or countable joins.

There is no assignment of values to its finitely many basic sentences, within their prescribed bounds and the intervals of $\mathfrak{p}_0$, that satisfies these ground conditions when terms are evaluated through their records. Indeed, (M3) and the positive antecedents give $\mu(u) = \mu(v)$ for each pair in $\mathcal{D}$. For each ordered basis, (M4) gives $\mu(d'_{i+1}) = \mu(d_i) + \mu(a'_i)$, and the recorded identities consequently give

$$\mu(d_{i+1}) = \mu(d_i) + \mu(a_i), \qquad \mu(a_F) = \sum_{i<4} \mu(a_i).$$

In the named basis of the block, the first summand is $\mu(\mathrm{at}(e_{n_0})) = \mu(h_{n_0}) > 15/16$ and hence is 1 by the dispersion-free condition and (M1). All other summands are nonnegative and $\mu(a_F) \leq 1$ by (M1), whence $\mu(a_F) = 1$. Each of the nine bases of Table 1 therefore has exactly one contextual atom term of measure 1. Since the two contextual terms of each shared ray have equal measures, the double count in the proof of Theorem C yields the same contradiction, between an odd and an even number of pairs.

Olivia now fairly challenges only the finitary commitments of $\mathfrak{p}_0 \cup \mathfrak{q}$ and their descendants under (Cn) and (Tm). All their closed terms already have records; (Tm) merely replaces them by their recorded witnesses. If this continuation survived indefinitely, the compact intervals for its finitely many basic sentences would have common values, because no position was infeasible. The induction on finitary formulas in the proof of Theorem 3.3, using all tolerances in the fair schedule, would make these values satisfy $\mathfrak{p}_0 \cup \mathfrak{q}$, contrary to the preceding paragraph. Thus the continuation wins. It makes no countable challenge and is finitely branching, so its tree is finite by Kőnig's lemma. Other commitments made earlier in the play may expose a contradiction sooner; they are not needed for this final argument.

Each of the preceding continuations ends after finitely many rounds in every play, and only finitely many of them follow the naming of N; the final parity continuation, moreover, has a finite tree. Once the enumerations are fixed and the least index is taken at each choice, these instructions determine Olivia's move at every finite play in which Josh has not yet lost. Every play in which she follows them is therefore finite, and the resulting strategy is a winning strategy with the two asserted phases. □

The strategy thus separates the countable search for localization and geometric margins from the final contradiction, which uses only the finite incidence pattern and finitary measure conditions. The parity configuration is pictured in Figure 8.

By Proposition 8.3(3), every winning strategy of Olivia in $\mathrm{G}(T^{01})$ has some play in which Josh gives a countable answer. For each fixed winning strategy and countable record, the lengths of plays with this record are bounded by Corollary 3.5(5); neither this bound nor the finite parity continuation gives a bound independent of the record. Whether such a bound can be achieved by a winning strategy remains open.

Question 8.7. Does Olivia have a winning strategy in $\mathrm{G}(T^{01})$ whose tree has finite rank, that is, such that every play in which she follows it ends within a fixed number of rounds?

Remark 8.8 (Standard models). The game $\mathrm{G}(T)$ concerns all models of T, standard or not, whereas Proposition 7.5 concerns the standard models, which are singled out by a hypothesis on the model rather than by axioms. For a Borel set B that meets $[0, 1]$, let T_B be the theory T together with the axiom $\sup_a \beta_B(a) \leq -1$, the normal form of the axiom of Proposition 7.5. Here the formula $\beta_B(a)$, given by Lemma 4.1, takes the value 1 at a when $\mu(a) \in B$ and -1 otherwise, in every structure that satisfies (Tr). By Theorem 3.3, Olivia

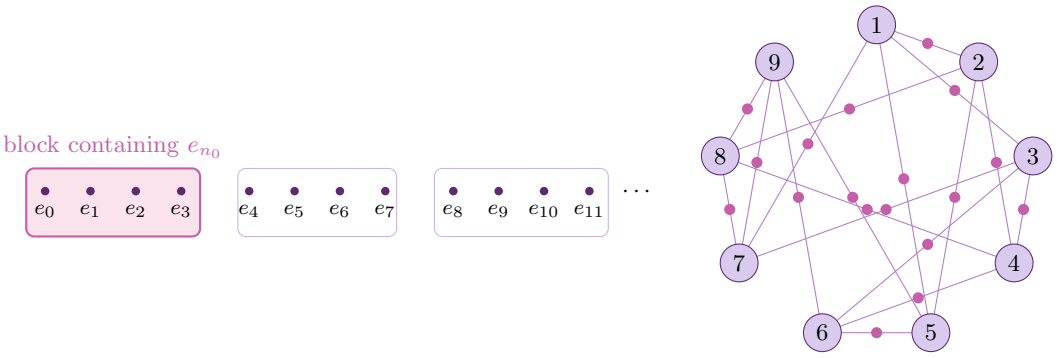


FIGURE 8. The parity configuration used in the final continuation of Proposition 8.6, shown schematically: the basis vectors e_n in blocks of four, with the block that contains e_{n_0} highlighted (here $n_0 = 2$), and the nine bases of Table 1 (numbered circles) with their eighteen rays (dots), each ray drawn on the segment that joins the two bases containing it.

has a winning strategy in $\mathrm{G}(T_B)$ exactly when T_B has no models at all, standard or not. This is the case when B contains 0 or 1, by Lemma 8.1(3) and (M2), and when B contains $(0, 1)$, by Corollary 8.2. Some other Borel sets B give theories with models. Indeed, the range of μ in a countable model $\mathfrak{N}$ of T (see Remark 10.3) is countable and contains 0 and 1 and, by Corollary 8.2, a point of $(0, 1)$. For the complement B of this range in $[0, 1]$, the structure $\mathfrak{N}$ is a model of T_B. We do not know which of the remaining Borel sets B give theories T_B with models.

9. Finite obstructions record what a Kochen–Specker configuration assumes

In the strategy of Proposition 8.6, countable additivity and noncontextuality enter at two distinct moments. Countable additivity selects the four-dimensional block in which the parity argument is carried out; within the block, the geometric axioms establish that the occurrences of a ray in different bases denote the same projection, and only then does axiom (M3) carry probabilities across these identities. We examine these two moments, after recalling what noncontextuality and determinism require of a hidden-variable model, and then give numerical tolerances under which the parity contradiction remains valid. As in Section 8.1, T^{01} denotes the theory T of Theorem B together with the requirement that μ take only the values 0 and 1.

Kochen and Specker certified, by a finite configuration of rays in a finite-dimensional Hilbert space, that no quantum probability measure on the closed subspaces of the space takes only the values 0 and 1. The bases of such a configuration are the sets of mutually orthogonal rays of the configuration that span the space. An assignment of the values 0 and 1 to its rays is *admissible* if it gives the value 1 to exactly one ray of each basis and to no two orthogonal rays. The configurations of Kochen and Specker have no admissible assignment.[6] We compare what such a configuration and the finite obstructions of the game record, how the measure enters each of them, and which topological argument gives their finiteness.

[6]An interactive presentation of the smallest Kochen–Specker configuration found so far in our ongoing computational search is maintained, together with the current state of this search, on the UTSA Quantum Mathematics website.

9.1. **Noncontextuality and determinism concern the responses at a single state.** One and the same projection can represent outcomes of different measurements: the projection onto a ray is an outcome of every measurement whose orthonormal basis contains a vector spanning the ray. To interpret the measure of the game, we must therefore keep three things apart: the outcome events themselves, the projections that represent them, and the conditional probabilities that a hidden-variable model assigns to them. Keeping them apart explains why, in the game, the geometric identities between projections and the requirements on the measure enter through separate axioms.

In the operational language of quantum foundations, an *outcome event* is a specified outcome of a specified measurement procedure. Quantum theory represents a projective measurement on a complex Hilbert space H by a family of mutually orthogonal projections that sum to the identity, which we call a *context*, and each of its outcome events by a member of this family. A preparation procedure is represented by a density operator ρ on H, a positive operator of trace 1, and the Born rule gives an event represented by a projection P the probability $\mathrm{Tr}(\rho P)$. Two outcome events are *operationally equivalent* when no preparation distinguishes them, that is, when their probabilities agree for every preparation; when all quantum preparations are available, this happens exactly when the two events are represented by the same projection.

A hidden-variable theory, in the form that Spekkens calls an *ontological model*, supposes that each preparation selects an underlying state according to a probability distribution, and that each underlying state determines, for every outcome event, a *response* in $[0,1]$, the conditional probability of the event given that state [Spe05]. At each underlying state, the responses to the outcomes of a measurement sum to 1, and averaging the responses over the distribution selected by a preparation must reproduce the operational probabilities. *Measurement noncontextuality* requires operationally equivalent events to receive equal responses at every underlying state, and *outcome determinism* requires the responses to take only the values 0 and 1. Both requirements concern the responses at individual underlying states, whereas operational equivalence concerns only their averages, which depend on the preparation.

Consider an ontological model that reproduces the quantum probabilities for all preparations and for the projective measurements on H with finitely many outcomes. At a fixed underlying state, measurement noncontextuality makes the response to an event depend only on the projection that represents it, and outcome determinism restricts the response to the values 0 and 1; together they yield a map from the projections of H to $\{0,1\}$. This map is normalized, since the responses within one measurement sum to 1, and it is finitely additive: for orthogonal projections P and Q, comparison of the measurements $\{P, Q, 1-P-Q\}$ and $\{P+Q, 1-P-Q\}$ shows that the response to $P+Q$ is the sum of the responses to P and to Q. On the whole lattice of projections, finite additivity already yields an obstruction: a fixed finite family of projections of infinite rank refutes every finitely additive $\{0,1\}$-valued measure (Section 9.2).

Countable additivity enters with the measurements that have infinitely many outcomes. A nondegenerate measurement in an orthonormal basis $(f_n)_{n\in\mathbb{N}}$ of an infinite-dimensional space H has as outcomes the rays of the vectors f_n, and the requirement that, at each underlying state, the responses to these outcomes sum to 1 is countable additivity along the basis. The predicate μ of T^{01} represents the map determined by the responses, and axiom (M5), which requires the measure to be continuous along the finite joins that exhaust a represented span, contains this normalization. The strategy of Section 8.1 uses (M5), together with the

instances of (K4), along the orthonormal sequence named by the constants e_n to locate a ray of measure close to 1 and the block of four consecutive basis vectors that contains the ray.

For a model that handles only measurements with finitely many outcomes, the obstruction needs projections of infinite rank and infinite corank: on the events of finite rank and their orthogonal complements, the finitely additive measure of Proposition 9.1 is a noncontextual valuation. With the nondegenerate measurements, countable additivity lets the rays carry the obstruction, and Section 9.2 shows how countable additivity chooses the rays.

The dispersion-free requirement thus concerns the responses at a single underlying state, whereas a density operator gives the Born probabilities of a preparation, which may take every value in $[0, 1]$.

9.2. **Countable additivity chooses the configuration.** A configuration of Kochen and Specker is fixed in advance, and one configuration serves against every assignment. Such a configuration can be realized in the laboratory, each basis of the configuration as a measurement; D'Ambrosio and his coauthors implemented the nine bases of Table 1 with single photons in 2013 [DHA+13]. Under the strategy of Section 8.1, the location of the ray configuration is determined by Josh's answers. In the localization phase of Proposition 8.6, the constants e_n, for $n \in \mathbb{N}$, give the orthonormal basis $(\hat{e}_n)_{n\in\mathbb{N}}$ of the Hilbert envelope $\bar{E}$, and (K4) prescribes the entries of ε by $\widehat{\kappa_n \varepsilon} = 2^{-n-1}\hat{e}_n$. During localization, a countable challenge to (M5) at the constant ε makes Josh name a number N and commit the measure of the join of the rays of $e_0, \ldots, e_{N-1}$ to a value greater than $7/8$. His later answers single out some $n_0 < N$ for which he commits the measure of the ray of e_{n_0} to a value greater than $15/16$. The eighteen rays then lie in the span F of the block $e_{4j}, \ldots, e_{4j+3}$ with $4j \leq n_0 < 4j + 4$.

Under this strategy, the countable-additivity scheme along the named orthonormal basis selects the four-dimensional block used for the parity argument, and the models of finite sets of axioms of T^{01} explain why this localization uses the entries prescribed for ε. In the models of the finite sets of axioms of T^{01} constructed in the proof of Proposition 8.3(2), the measure vanishes on every finite-dimensional subspace. These structures satisfy (M5) on the sequences present in their restricted sequence sorts. Their measures are nevertheless not σ-additive on the lattice L_{fc} of Example 6.5: the rays of the vectors e_n are mutually orthogonal and have join 1 in this lattice, yet each of them has measure 0. Only σ-additivity over all the vectors e_n, through infinitely many instances of (K4), excludes a measure vanishing on every finite-dimensional subspace. In finite dimension this step is automatic, since a configuration of Kochen and Specker spans the space.

The models of the finite sets of axioms of T^{01} described above rest on the following elementary fact, which isolates what finite additivity cannot exclude on finite-dimensional events.

Proposition 9.1. *Let H be a separable infinite-dimensional complex Hilbert space, and let $L_{\mathrm{fc}}(H)$ be the set of finite-dimensional subspaces of H and their orthogonal complements. Then $L_{\mathrm{fc}}(H)$ is a sub-ortholattice of $L(H)$, and the map μ_{fc} that vanishes on the finite-dimensional subspaces and takes the value 1 on their orthogonal complements is a finitely additive $\{0, 1\}$-valued probability measure on $L_{\mathrm{fc}}(H)$: $\mu_{\mathrm{fc}}(H) = 1$, and $\mu_{\mathrm{fc}}(a \vee b) = \mu_{\mathrm{fc}}(a) + \mu_{\mathrm{fc}}(b)$ whenever a and b are orthogonal elements of $L_{\mathrm{fc}}(H)$.*

Proof. Up to unitary equivalence $H = \ell^2$, and $L_{\mathrm{fc}}(H)$ is then the sub-ortholattice L_{fc} of $L(\ell^2)$ of Example 6.5. Since H is infinite-dimensional, no subspace is both finite-dimensional and the orthogonal complement of a finite-dimensional subspace, so the map μ_{fc} is well defined, and $\mu_{\mathrm{fc}}(H) = 1$ because $H = \{0\}^{\perp}$. Let a and b be orthogonal elements of $L_{\mathrm{fc}}(H)$. If $a = G^{\perp}$ with G finite-dimensional, then $b \subseteq a^{\perp} = G$, so b is finite-dimensional. Hence at most one of a and b is infinite-dimensional, and $a \vee b$ is infinite-dimensional exactly when one

of them is. Both sides of the additivity equation therefore equal 1 when one of a and b is infinite-dimensional, and 0 otherwise. □

A configuration of rays inside a finite-dimensional subspace F can refute a $\{0,1\}$-valued measure only when the measure gives F the value 1, and μ_{fc} gives every finite-dimensional subspace the value 0.

On the full projection lattice, finite additivity alone already gives an obstruction using infinite-rank projections. Write a separable infinite-dimensional Hilbert space H as $\mathbb{C}^4 \otimes \ell^2$, and, for each vector v of Table 1, let P_v be the projection of $\mathbb{C}^4$ onto the ray of v. For each basis $\mathcal{B}$ of the table, the four projections $P_v \otimes 1$ with v in $\mathcal{B}$ are pairwise orthogonal with sum 1. The eighteen projections $P_v \otimes 1$, which have infinite rank, therefore form a configuration that is fixed in advance and refutes every finitely additive $\{0,1\}$-valued probability measure on $L(H)$ by the same parity argument. More generally, Döring proved the Kochen–Specker theorem for von Neumann algebras without summands of types I_1 and I_2 [Dör05]. His argument passes through the finitely additive form of Gleason's theorem due to Christensen, Yeadon, and Bunce and Wright [Chr82, Yea83, BW92] (see Hamhalter [Ham03]).

The Solèr sort of the models of finite sets of axioms of T^{01} constructed in the proof of Proposition 8.3(2) is the lattice L_{fc} of Example 6.5. This lattice contains no subspace of infinite dimension and infinite codimension, and in particular none of the ranges of the projections $P_v \otimes 1$. The $\{0,1\}$-valued measures of these models therefore do not contradict the finite-additivity obstruction on the full lattice $L(H)$.

A configuration of Kochen and Specker certifies the impossibility by the absence of an admissible assignment, whereas a finite obstruction certifies it by a finite position of commitments and function records that no structure satisfies. The first kind of certificate rests on compactness: as recalled in the Introduction, when the Hilbert space is finite-dimensional, the condition that an orthonormal basis contain exactly one ray of value 1 defines a closed set of assignments, and finitely many of these closed sets already have empty intersection, which turns the theorem of J. S. Bell into a finite configuration [Bel96, Pit98]. The condition that the probabilities of the rays of the vectors e_n add up to 1 defines a set that is not closed, and conditions of this kind, expressed by the countable suprema and infima of real-valued $L_{\omega_1\omega}$, lie beyond the reach of compactness. For this logic, the space of Henkin types is Čech-complete but in general not compact, and the Baire category theorem in that space supports the omitting types argument recalled in Section B. The finite obstructions of the game come from model construction along fair plays and the determinacy theorem of Gale and Stewart; the bookkeeping of this construction has its own Baire category reading, in the space of plays, which Section B sets beside the argument in the space of Henkin types.

9.3. **The game records noncontextuality as the substitutivity of equality.** We continue with the plays of the strategy of Section 8.1 that reach the parity argument. In a usual Kochen–Specker configuration, a ray is one vertex in all the bases that contain it, so one value for the ray across these bases is part of the data. A *contextual vertex*, by contrast, indexes an occurrence of a term representing a ray in a specified basis; each ray of the eighteen-ray configuration has two such vertices. In every model of T, the two terms denote elements a and b with $[a = b] = 1$, meaning that the corresponding orthogonal projections $\bar{P}_a$ and $\bar{P}_b$ of the Hilbert envelope $\bar{E}$ are equal (see Lemma 5.3 and part (3) of Lemma 5.4). The elements themselves may be distinct: Keisler's general structures, on which the logic rests, have no equality symbol [Kei23, Section 2.1], and even a standard model may contain distinct elements representing the same projection. The finite position in Proposition 8.6 records positive lower bounds on the formulas expressing these identities. The corresponding

ground instances of (M3), with closed terms substituted for the variables, then force equality of the measure values; a positive antecedent suffices, without requiring the value 1. Figure 9 illustrates this transfer between contextual vertices.

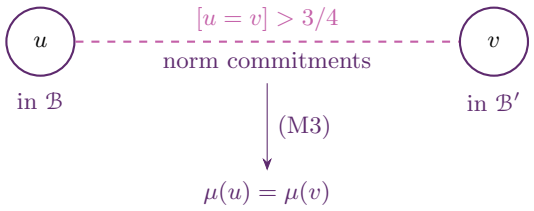


FIGURE 9. Probability transport across a shared ray in Proposition 8.6, between terms u and v that represent the same ray in the bases $\mathcal{B}$ and $\mathcal{B}'$.

Representing contextual occurrences separately has precedents in quantum foundations. For the eighteen-ray configuration, Kunjwal and Spekkens [KS15, Fig. 2] represent measurement events by distinct nodes and display the operational equivalences between them. Our analysis places the relevant projection identities within a theory of vectors, sequences and subspaces, and follows their consequences through the probability predicate and the model existence game.

Other frameworks distinguish contextual occurrences through different mathematical requirements. In the contextuality-by-default approach of Dzhafarov and Kujala, random variables with a common content in different contexts are represented separately. A coupling is a jointly distributed family reproducing the joint distribution in each context; the noncontextuality criterion imposes specified maximal-agreement requirements on occurrences with common content [KDL15, DKL15]. In the sheaf-theoretic framework of Abramsky and Brandenburger, an empirical model assigns to each context a probability distribution on the joint outcomes of its measurements, and these distributions agree on the overlaps of contexts. Noncontextuality requires a probability distribution on assignments of outcomes to all measurements whose marginal on each context reproduces the distribution assigned to this context [AB11]. The present comparison concerns the representation of cross-context relations: our certificates track the derivation of projection identities and their use by a dispersion-free valuation.

Table 2 compares the finite obstructions with the usual ray configuration, separating what holds for every winning strategy of Olivia from what belongs to the strategy of Section 8.1.

The geometric identifications in Table 2 come from the terms used to construct each basis. As in the proof of Theorem C, each of the nine bases $\mathcal{B}$ of Table 1, with vectors $v_0, \dots, v_3$, gives a closed term $s_{\mathcal{B}}$ of the sequence sort whose entries $\kappa_i s_{\mathcal{B}}$ satisfy $\widehat{\kappa_i s_{\mathcal{B}}} = \frac{1}{16} v_i$ for $i < 4$. Since the subspace $S_{\mathrm{at}(x)}$ is the ray $\mathbb{C}\hat{x}$ whenever $\|x\| > 0$ (Lemma 5.4), the vertices of $\mathcal{B}$ are indexed by the terms $\mathrm{at}(\kappa_i s_{\mathcal{B}})$, for $i < 4$. Every ray of the table lies in two bases and is therefore represented by two contextual occurrences of terms.

The instances of (K2), together with those of the axioms behind Lemmas 5.3 and 5.4, give $[\mathrm{at}(\kappa_i s_{\mathcal{B}}) = \mathrm{at}(v_i)] = 1$, where $\mathrm{at}(v_i)$ is as in the proof of Theorem C; hence the two contextual terms of a ray denote elements that are equal in the sense of the formula $[a = b]$. The same instances give $[c_{\mathcal{B}} = a_F] = 1$ for the nine elements $c_{\mathcal{B}} = \mathrm{at}(\kappa_0 s_{\mathcal{B}}) \vee \cdots \vee \mathrm{at}(\kappa_3 s_{\mathcal{B}})$, where the element a_F stands for the span F of the block. All these identifications follow from axioms that do not mention μ.

	Kochen–Specker configurations	finite obstructions of the model existence game for T^{01}
framework	partial Boolean algebras: in dimension at least three, the projections admit no homomorphism into the two-element Boolean algebra	real-valued $L_{\omega_1\omega}$: the theory T^{01} has no models, although each of its finite sets of axioms has one
		for every winning strategy of Olivia, who claims that T^{01} has no models
tree of plays	none: one configuration serves against every assignment	well-founded, and branching infinitely at some node (Corollary 8.4)
certificate of impossibility	no assignment of 0 and 1 to the rays gives 1 to exactly one ray of each basis and to no two orthogonal rays	at each leaf, a finite position of commitments and function records satisfied in no structure
		under the strategy of Section 8.1, in a play that lasts until Olivia has carried out the parity argument of Theorem C
bases	orthonormal bases of the space	the nine bases of Table 1 in the span F of a block, each completed by $F^\perp$
vertex	a ray	a contextual occurrence of a term representing a ray in one of the nine bases
identifications	none needed: a ray is one vertex in all the bases that contain it (the hypothesis of noncontextuality)	principally eighteen identifications of rays and nine of the joins of the bases with F, all derived from axioms of T that do not mention μ, along which axiom (M3) carries probabilities; the parity argument needs all eighteen identifications of rays
values	0 or 1 at each ray	0 or 1 for the probability of each contextual occurrence, imposed by ground instances of (M1) and of the dispersion-free axiom
constraint on a basis	exactly one ray of value 1	sum 1 of the probabilities of its occurrences, imposed by ground instances of (M1), (M3) and (M4) once an atom of the block has a recorded value above 15/16; by Theorem 9.2, the contradiction persists under tolerances δ on these values and η on these sums with $4\delta + \eta < 1$
location of the rays	fixed in advance	the eighteen rays of Table 1, in the block determined by Josh's answers
space	finite dimension, spanned by the rays	infinite dimension; the rays lie in the span of a block $e_{4j}, \dots, e_{4j+3}$ of the orthonormal basis $(e_n)_{n\in\mathbb{N}}$ named by constants of T
role of σ-additivity	none: every basis is finite; in infinite dimension, the eighteen projections $P_v \otimes 1$ on $\mathbb{C}^4 \otimes \ell^2$, described in Section 9.2, need only finite additivity	locates the unit-weight block: Josh names a number N and commits the measure of the join of the rays of $e_0, \dots, e_{N-1}$ to a value greater than 7/8, and the eighteen rays lie in the block that contains some e_{n_0} with $n_0 < N$

TABLE 2. The certificates of the model existence game for T^{01} beside the finite coloring certificate of Kochen and Specker: first for every winning strategy of Olivia, then under the strategy of Section 8.1, for plays that last until Olivia has carried out the parity argument of Theorem C.

Axiom (M3), $[a = b] \wedge |\mu(a) - \mu(b)| \leq 0$, which every quantum probability measure satisfies, says that μ respects the equality of the Solèr sort. This axiom requires the measure to take the same value at the elements denoted by the two contextual terms of a ray. Noncontextuality, which a configuration of Kochen and Specker builds into the identification of a ray across the bases that contain the ray, therefore enters the certificate exactly as these ground instances of (M3).

For the following minimality assertion, retain only the Boolean variables attached to the thirty-six contextual occurrences, the nine exactly-one basis constraints and the eighteen equalities between paired occurrences. Each equality is indispensable to the inconsistency of this constraint system. Indeed, the nine bases can be arranged in the cycle $1, 2, 4, 3, 6, 5, 9, 8, 7$ of their numbers in Table 1, in which any two consecutive bases (the last and the first included) share a ray. If the identification attached to a ray of a basis $\mathcal{B}$ is omitted, give the value 1 to the vertex of the ray in $\mathcal{B}$. Next, pair off the other eight bases, which follow one another along the cycle, by four of the rays that consecutive bases share. Finally, give the value 1 to both vertices of each of these four rays and the value 0 to all the other vertices. Each basis then has exactly one vertex of value 1. Without identifications, a vertex of value 1 could be chosen in each basis independently. An independent choice is possible in the dense families of orthonormal bases that Clifton and Kent constructed in each finite dimension. In each of these families, any two bases generate Boolean algebras of projections with only 0 and the identity in common [CK00].

9.4. **The parity obstruction survives numerical tolerances.** The interval commitments of the game admit a quantitative stability statement for the parity argument. For a hypothetical valuation, we attach a value to each occurrence of a ray in a basis and allow these values to be near 0 or 1, with each basis sum near 1. The following theorem gives sufficient numerical tolerances while keeping the rays fixed and the values at the two occurrences of a shared ray exactly equal.

Theorem 9.2 (Quantitative parity obstruction). *Let $\mathfrak{B}$ be the family of the nine bases of Table 1, each regarded as a set of four of the eighteen rays determined by the table. Let δ and η be nonnegative real numbers with $4\delta + \eta < 1$. There is no family of numbers $(p_{B,r})_{B \in \mathfrak{B},\, r \in B}$ in $[0, 1]$ satisfying all three conditions:*

(i) $p_{B,r} = p_{B',r}$ whenever $B, B' \in \mathfrak{B}$ and $r \in B \cap B'$;
(ii) $p_{B,r} \in [0, \delta] \cup [1 - \delta, 1]$ for every $B \in \mathfrak{B}$ and every $r \in B$;
(iii) $\left|\sum_{r \in B} p_{B,r} - 1\right| \leq \eta$ for every $B \in \mathfrak{B}$.

Proof. Since $4\delta + \eta < 1$, the intervals $[0, \delta]$ and $[1 - \delta, 1]$ are disjoint. Call an occurrence, that is, a pair (B, r) with $B \in \mathfrak{B}$ and $r \in B$, high when $p_{B,r}$ belongs to $[1 - \delta, 1]$. A basis with no high occurrence has sum at most $4\delta < 1 - \eta$, contrary to (iii). A basis with at least two high occurrences has sum at least $2 - 2\delta > 1 + \eta$, again contrary to (iii). Thus each of the nine bases has exactly one high occurrence. Condition (i) gives both occurrences of each ray the same value, and each ray belongs to exactly two bases. Counting high occurrences basis by basis therefore gives nine, whereas counting them ray by ray gives an even number, a contradiction. □

The bound $4\delta + \eta < 1$ is sharp for these three conditions. If $4\delta + \eta \geq 1$ and $\eta \leq 1$, the family with every $p_{B,r}$ equal to $(1 - \eta)/4$ satisfies them, since $(1 - \eta)/4 \leq \delta$ and every basis sum is $1 - \eta$; if $\eta \geq 1$, the family that vanishes identically satisfies them. The sharpness

concerns the numerical model of Theorem 9.2, in which the rays and the equalities across bases are exact, and not the noise tolerance of an experiment.

In the finite parity continuation of Proposition 8.6, the three conditions of the theorem appear with $\delta = \eta = 0$. Condition (i) comes from the ground instances of (M3) and the recorded geometric antecedents. The dispersion-free ground conditions, together with (M1), give condition (ii). Localization has recorded an atom value greater than 15/16, which these conditions raise to 1; finite additivity, nonnegativity and the upper bound on the block measure then force each of the nine basis sums to be 1, which is condition (iii).

The quantitative theorem isolates the same parity obstruction when endpoint and basis-sum conditions are allowed numerical tolerances: for bases of k rays, with an integer $k \geq 2$, the bound $k\delta + \eta < 1$ forces exactly one high value per basis. Any configuration whose exactly-one constraints have no Boolean solution then gives the same contradiction.

The parameters δ and η bound errors in commitments about a hypothetical dispersion-free valuation. Born probabilities may take intermediate values, so condition (ii) is a restriction on this valuation rather than an estimate of experimental frequency errors. An experimental noncontextuality inequality would additionally require specified preparation and measurement procedures, operational equivalences and a statistical bound for the relevant class of ontological models. Theorem 9.2 thus establishes the quantitative stability of the obstruction to a dispersion-free valuation when the rays are exact; the distinct problem of imperfect measurements has its own analyses [SBZ01, Lar02].

Historical note. In 1990 Peres gave a proof of the Kochen–Specker theorem by observables on two qubits, valid for a particular state. Building on Peres's proof, Mermin found a finite obstruction of another kind, formed by nine such observables arranged in three rows and three columns of commuting observables. No assignment of the values ± 1 to these nine observables respects the multiplicative relations along the rows and the columns [Per90, Mer90]. The array thus gives a state-independent proof of the theorem. In 1999 Meyer observed that a coloring of Godsil and Zaks [GZ88] assigns 0 and 1 to the unit vectors of $\mathbb{R}^3$ with rational coordinates, which are dense in the unit sphere. In this assignment, each orthogonal triple contains exactly one vector of value 1, and no two orthogonal vectors have the value 1 [Mey99]. Kent extended the result in the same year to dense sets of resolutions of the identity in every finite dimension [Ken99], and the families of Clifton and Kent date from 2000 [CK00].

10. Remarks

The vocabulary and axioms determine how the theory describes subspaces and their measures, and which representation results require standardness. The following remarks explain the use of projections, the role of (K4) in Theorem C, and the choices of scalars and logical operations. They also distinguish the conclusions proper to standard models from the expressive limits of the axioms and record two redundancies among the axioms. None is needed for the proofs of the main theorems.

Remark 10.1 (Projections rather than membership). The join of two closed subspaces is their closed span, so a description of the join must account for closure. The vocabulary does this through projections. A vector x belongs to an element a of the Solèr sort exactly when $\|\frac{1}{2}x - \frac{1}{2}\pi(a, x)\| = 0$. One might instead try to describe membership in $a \vee b$ by asking whether x is a sum $y + z$, with y in a and z in b. In the logic, this would mean that the infimum of a scaled norm of $x - y - z$ vanishes. The vectors quantified over lie in the unit ball, and this restriction prevents the description from recovering the closed span. Consider a standard model whose Hilbert sort is the closed unit ball of a Hilbert space H. In this model,

the infimum of the scaled norm of $x-y-z$ vanishes exactly at the sums $y+z$ of such members, since these sums form a closed set. Indeed, bounded sequences in the closed subspaces S_a and S_b spanned by the members of a and of b have weakly convergent subsequences. When $S_a + S_b$ is not closed, these sums therefore miss the vectors of the unit ball that lie in the join $\overline{S_a + S_b}$ of S_a and S_b in $L(H)$ but not in $S_a + S_b$. Axiom (Sub2) reaches the join through the projections instead, by von Neumann's alternating projection theorem. The atoms are reached in the same manner, since axiom (At1) describes $\mathrm{at}(x)$ through the projection onto it: in a standard model of T_0, the functions at and rep provide the bijection between the rays of H, which correspond to the unit vectors up to a phase, and the atoms of the quotient of the Solèr sort by its equality (see Proposition 6.3). Completeness of this quotient, by contrast, is imposed by no axiom: in the standard models of T_1 it is a theorem, which the sequence sort yields through the function span (see Theorem 7.3).

Remark 10.2 (Finitary definability). Every closed interval with rational endpoints is definable over T by a formula with countable joins and without quantifiers (see the proof of Lemma 4.1), and with quantifiers the definable sets are exactly the σ-projective sets. We do not know whether a formula without countable joins can define over T any subset of $[0,1]$ other than $\emptyset$ and $[0,1]$. An argument through ultraproducts and the continuity of formulas would have to overcome two features of the standard models of T. An ultraproduct of standard models of T need not satisfy axiom (H2), which places every vector in the closed span of the vectors e_n; and the lattice operations are not continuous for the distance between orthogonal projections in operator norm, since the meet of two distinct rays is the zero subspace however close the rays are.

Remark 10.3 (Nonstandard models). The representation theorems identify the lattice and the measure in standard models. Their conclusions cannot extend to every model of T: the downward Löwenheim–Skolem theorem gives models whose sorts are countable [ADI26, Lemma 4]. In such a model, the range of μ is countable and hence differs from $[0,1]$. The quotient of the Solèr sort by $[a = b] = 1$ is countable as well, so it cannot be isomorphic to the uncountable lattice of closed subspaces of an infinite-dimensional Hilbert space. Thus standardness is essential to Proposition 7.5 and to part (1) of Theorem B.

In a nonstandard model of T_0, even the conclusion of Gleason's theorem for T_0 (see Theorem 6.6) can fail, as the pre-Solèr lattice of Example A.21 shows. Recall from this example that E_0 is the space of finitely supported complex sequences with the inner product of ℓ^2, and that Λ_0 is the lattice formed by the finite-dimensional subspaces of E_0 and their orthogonal complements in E_0. Take the closed unit ball of E_0 as Hilbert sort, the standard unit vectors as the constants e_n for $n \in \mathbb{N}$, and the lattice Λ_0 as Solèr sort. The function π sends (a, x) to the orthogonal projection of x onto a, and the function at sends x to $\mathbb{C}x$. Finally, the function rep assigns to each $a \neq \{0\}$ a unit vector $\mathrm{rep}(a)$ chosen in a, and to $\{0\}$ the vector 0. With the truncations interpreted by their intended values, the resulting structure is a model of T_0 that is not standard, since E_0 is not complete. Passage to closures in ℓ^2 embeds the ortholattice Λ_0 in the lattice L_{fc} of Example 6.5, and each map $\pi(a, \cdot)$ is the restriction to the unit ball of E_0 of the orthogonal projection P_a of ℓ^2 onto the closure of a; axioms (Sub1) and (Sub2) therefore hold by Theorem 5.1, applied in ℓ^2, and the other axioms of T_0 hold by construction. Every maximal orthonormal family of E_0 is countable, and its rays have join 1 in Λ_0. A σ-additive state ν on Λ_0 would therefore give a function f on the unit vectors of E_0, defined by $f(x) = \nu(\mathbb{C}x)$, with values in $[0,1]$. Its sum over every maximal orthonormal family would be 1, and its sum over an orthonormal basis of any finite-dimensional subspace M would be $\nu(M)$, independently of the basis. Dvurečenskij calls a nonnegative function on the unit vectors of an inner product space a *weak frame function* when its sums over

all maximal orthonormal families equal one positive constant and its restriction to the unit vectors of every finite-dimensional subspace has the same sum over all the orthonormal bases of the subspace [Dvu95]. The function f would therefore be a weak frame function, and by his theorem an inner product space that carries a weak frame function is complete [Dvu95, Theorem 3.1] (compare Hamhalter and Pták [HP87]). The space E_0, however, is not complete. Hence the lattice Λ_0 has no σ-additive state, although every density operator ρ on ℓ^2 defines a map $a \mapsto \operatorname{Tr}(\rho P_a)$ on Λ_0. The identification of σ-additive states with all such trace maps therefore fails in this nonstandard model.

In standard models, the axioms of Proposition 7.5 exclude every range other than $[0,1]$, one Borel set at a time, or more generally one σ-projective set at a time (see Theorem 7.7). No sentence, however, and indeed no countable theory, says that the range of μ is disconnected. Indeed, suppose that a countable theory T' of real-valued $L_{\omega_1\omega}$ in the vocabulary of T were satisfied, among the models of T, exactly by those in which the range of μ is disconnected; a sentence ψ that took the value 1 exactly in these models would give such a theory, namely the closed condition $\psi = 1$. Let $\mathfrak{M}$ be a standard model of T, which exists by Theorem B(2). By Proposition 7.5, the range of μ in $\mathfrak{M}$ is the connected set $[0,1]$, which excludes $\mathfrak{M}$ from the models of T'. By the downward Löwenheim–Skolem theorem [ADI26, Lemma 4 and its proof], applied to a countable fragment that contains the sentences of the axioms of T and of the conditions of T', the model $\mathfrak{M}$ has a substructure $\mathfrak{N}$ with countable sorts in which every sentence of the fragment keeps its value. The substructure $\mathfrak{N}$ is therefore a model of T but not of T'. The range of μ in $\mathfrak{N}$, however, is countable and contains $\mu(0) = 0$ and $\mu(1) = 1$; it is therefore disconnected, which contradicts the choice of T'.

We leave open whether the measure μ of a nonstandard model of T can differ from every map $a \mapsto \operatorname{Tr}(\rho \bar{P}_a)$ with ρ a density operator on a complex Hilbert space $\bar{E}$ as in Lemma 5.3. Here $\bar{P}_a$ is the orthogonal projection of $\bar{E}$ that extends $\pi(a,\cdot)$.

Remark 10.4 (Axiom (K4)). Unlike the representation theorems, Theorem C excludes a dispersion-free measure in every model of T, standard or not. Axiom (K4) prescribes the entries of the constant ε by $\widehat{\kappa_n \varepsilon} = 2^{-n-1}\hat{e}_n$, for $n \in \mathbb{N}$; together with (H2) and (Sp1), these identities make the subspace represented by $\operatorname{span}(\varepsilon)$ the whole Hilbert envelope (see Lemma 8.1(4)). Without (K4), a nonstandard model of the remaining axioms may contain no sequence whose entries span its Hilbert envelope. The structure in the proof of Proposition 8.3(2), with the zero sequence as ε, satisfies every other axiom of T and has a measure taking only the values 0 and 1. In standard models, Gleason's theorem already excludes such a measure by Theorem B, whose proof does not use (K4).

Remark 10.5 (Scalars). The vocabulary fixes the complex numbers as scalars. For real Hilbert spaces, replace $\mathbb{Q}(i)$ by $\mathbb{Q}$ and omit the predicate $\operatorname{Im}\langle x \mid y\rangle$. For quaternionic spaces, use the quaternions with rational coordinates, acting on the left, and four predicates for the real components of the form. In Proposition 6.3, "complex number of modulus one" then becomes "real number of modulus one" or "quaternion of modulus one". Each choice yields the analogue of Theorem 7.3 for the corresponding case of Theorem A.18, and the proof of Theorem C applies in all three cases, since the vectors of Table 1 have real coordinates.

Remark 10.6 (The logical form of the axioms). Every axiom of T is a closed condition of real-valued infinitary logic, built with the connectives $s \cdot \varphi$, $\varphi + \psi$, $\varphi \vee \psi$, the suprema over the sorts and countable joins of families with bounded a priori bounds. The axioms (H1), (OL), (P), (Sub0), (Sub3), (At1), (At2), (K1), (K2), (K4), (Sp1), (Tr) and (M1)–(M4) involve no countable join, except through the squares of Lemma 2.1 in (H1), (K1) and (At1). Axioms (H2), (Sub1), (Sub2), (K3), (Sp2) and (M5) use countable suprema or infima, and (At3) uses the countable join $\lceil \|x\| > 0 \rceil$. The hypothesis of the implication in (M3) is the

formula $[a = b]$ of Section 4, and that of (Sp2) is the countable conjunction of the formulas $[\mathrm{at}(\kappa_n s) \leq b]$. These formulas are built from the norms of the vectors $\mathrm{rep}(c)$, and they are two-valued in every model of T_1, by Lemma 5.4 applied to its reduct to the vocabulary of T_0; hence so is their conjunction. In every such model they express the order and the equality of the Solèr sort. Each axiom (Tr), for given R, n and q, consists of closed conditions on formulas without countable joins, whose a priori bounds grow with n; the formulas $\lceil R(\bar{t}) > q \rceil$ and $\lceil R(\bar{t}) < q \rceil$, which gather the truncations for all n, are countable joins of members with a priori bound 1. The functions at and span take values in the Solèr sort, on which the formula $d(a, b) = \frac{1}{2}(1 - [a = b])$ defines, in every model of T, a pseudometric with the values 0 and 1. Its topology induces the discrete topology on the quotient by the equality. On the subspace of rays $\mathbb{C}\hat{x}$ with $\hat{x} \neq 0$, represented by the classes of $\mathrm{at}(x)$, this topology differs from the strong operator topology that Section 7.1 places on the projections. In a standard model, the function at sends arbitrarily close nonzero, nonproportional vectors to distinct quotient classes, so it is not continuous from the norm topology to the discrete quotient topology. The function span has the same discontinuity along the sequences $\iota_0 x$ and $\iota_0 y$, since Lemma 7.2 identifies the class of $\mathrm{span}(\iota_0 x)$ with the class of $\mathrm{at}(x)$. The logic accommodates this discontinuity, since it places no condition on the interpretation of a function symbol beyond its sorts. This observation concerns the norm topology and the discrete quotient topology just specified. It does not preclude a pre-metric expansion in Keisler's sense: for a countable vocabulary such an expansion exists with a distance predicate that is a uniform limit of formulas [Kei23, Theorem 3.3.4], and its topology need not be the norm topology.

Remark 10.7 (Redundancy). The axioms (OL) hold in every structure that satisfies the other axioms of T_0, since the proof of Lemma 5.4, which embeds the quotient of the Solèr sort into the lattice of closed subspaces of a Hilbert space, does not use them. In standard models, axiom (Sub2) likewise follows from (Sub1), (Sub3) and (OL). The two arguments do not allow both groups to be omitted together, because the proof of Lemma 5.4 uses (Sub2) to compute joins, whereas the derivation of (Sub2) uses (OL); both groups of axioms are kept nonetheless, so that each property has an axiom of its own.

Historical note. Löwenheim proved in 1915 that a first-order sentence with a model has a countable model [Lö15]. Skolem extended the result to countable sets of sentences in 1920, and in 1922 he observed that axiomatic set theory, in which uncountable sets exist, therefore has countable models [Sko70]. The countable models of Remark 10.3 are of the same kind: they satisfy every axiom of a theory whose standard models are uncountable. That an incomplete inner product space can be recognized by the absence of states goes back to Hamhalter and Pták. In 1987 they proved that a separable inner product space is complete exactly when the lattice of its orthogonally closed subspaces admits a σ-additive state [HP87]. The theorem of Dvurečenskij of 1995 used in Remark 10.3, which detects completeness through frame functions rather than through states on the whole lattice of orthogonally closed subspaces, belongs to the same line of work [Dvu95].

Appendix A. Lattices, hermitian spaces and projective geometry

In this appendix we give the geometric definitions and representation results used in Sections 6 and 7, together with their proofs: the MacNeille completion, hermitian spaces, coordinatization, Solèr's condition and the completion of a pre-Solèr lattice. These results describe the abstract lattices represented by the vector and event axioms; the model existence game uses only the logic of Section 2.

A.1. **Lattices.** Birkhoff and von Neumann proposed in 1936 that the propositions about a quantum system be identified with the closed subspaces of its Hilbert space, ordered by inclusion [BvN36]; the order-theoretic notions below describe this geometry without reference to vectors. The completion that MacNeille constructed in 1937 [Mac37] supplies the joins and meets that an incomplete lattice of subspaces lacks; the lattices that the axioms on vectors and events alone describe may be incomplete in this way.

Definition A.1. A *partially ordered set* is a set with a reflexive, antisymmetric and transitive relation $\leq$. For a subset S of a partially ordered set, an *upper bound* of S is an element b with $a \leq b$ for all $a \in S$, and the *join* $\bigvee S$ of S is its least upper bound, when it exists; lower bounds and the *meet* $\bigwedge S$ are defined dually. A *lattice* is a partially ordered set in which every pair of elements a, b has a join $a \vee b$ and a meet $a \wedge b$. A lattice is *bounded* when it has a least element 0 and a greatest element 1, and *complete* when every subset has a join.

Definition A.2. An *orthocomplementation* on a bounded lattice Λ is a map $a \mapsto a^\perp$ from Λ to Λ such that, for all $a, b \in \Lambda$, the equations $a^{\perp\perp} = a$, $a \wedge a^\perp = 0$ and $a \vee a^\perp = 1$ hold, and the inequality $a \leq b$ implies $b^\perp \leq a^\perp$. A bounded lattice with an orthocomplementation is an *ortholattice*. Two elements a, b of an ortholattice are *orthogonal*, written $a \perp b$, when $a \leq b^\perp$. An ortholattice is *orthomodular* when, for all elements a and b, the inequality $a \leq b$ implies $b = a \vee (b \wedge a^\perp)$.

The lattice $L(H)$ of closed subspaces of a Hilbert space H, ordered by inclusion and equipped with the orthocomplementation $M \mapsto M^\perp$, is an ortholattice. Its orthomodularity says that every closed subspace N containing a closed subspace M is the orthogonal sum of M and $N \cap M^\perp$.

Definition A.3. An *atom* of a bounded lattice is a nonzero element p such that every element b with $b \leq p$ is equal to 0 or to p. A bounded lattice is *atomistic* when every element a is the join of the atoms below it (this join being required to exist). It has the *covering property* when, for every atom p and every element a with $p \not\leq a$, no element lies strictly between a and $a \vee p$. It is *irreducible* when, for any two distinct atoms p and q, some atom r distinct from p and q lies below $p \vee q$.

Definition A.4. A *sub-ortholattice* of an ortholattice Λ is a subset that contains 0 and 1 and is closed under the operations $\vee$, $\wedge$ and $^\perp$ of Λ. An *embedding* of ortholattices is an injective map that preserves $\vee$, $\wedge$, $^\perp$, 0 and 1.

Let Γ be a complete lattice, and let Λ be a subset of Γ. We say that Λ is *join-dense* in Γ when every $g \in \Gamma$ satisfies $g = \bigvee\{a \in \Lambda : a \leq g\}$, the join being computed in Γ, and that Λ is *meet-dense* in Γ when every $g \in \Gamma$ satisfies $g = \bigwedge\{a \in \Lambda : a \geq g\}$. For instance, the set $\mathbb{Q}$ of rational numbers is join-dense and meet-dense in the extended real line $[-\infty, \infty]$.

For a partially ordered set Λ and a subset X of Λ, let X^u and X^l be the sets of upper and lower bounds of X in Λ. We call X a *cut* when $X = (X^u)^l$. Ordered by inclusion, the cuts form a complete lattice $\mathrm{DM}(\Lambda)$, the *MacNeille completion* of Λ. The map $a \mapsto \{b \in \Lambda : b \leq a\}$ embeds Λ into this lattice, preserving all joins and meets that exist in Λ.

When Λ is an ortholattice, its orthocomplementation induces one on the cuts. For a cut X, let $X^\perp$ be the set $\{b \in \Lambda : b \leq a^\perp$ for every $a \in X\}$ of lower bounds of $\{a^\perp : a \in X\}$. We show after the next proposition that the map $X \mapsto X^\perp$ is an orthocomplementation of $\mathrm{DM}(\Lambda)$ that extends that of Λ; with it we regard $\mathrm{DM}(\Lambda)$ as an ortholattice.

Proposition A.5 (Banaschewski [Ban56], Schmidt [Sch56]). *Let Λ be a partially ordered set, and let Γ be a complete lattice that contains Λ as a join-dense and meet-dense subset, the order of Γ inducing on Λ the order of Λ. Then the map that sends each $g \in \Gamma$ to the set $\{a \in \Lambda : a \leq g\}$ is an isomorphism of Γ onto $\mathrm{DM}(\Lambda)$ that restricts on Λ to the canonical*

embedding $a \mapsto \{b \in \Lambda : b \leq a\}$. *If* Λ *is an ortholattice, then the lattice* Γ *has at most one orthocomplementation that extends that of* Λ*; any such orthocomplementation is given by* $g^\perp = \bigwedge\{a^\perp : a \in \Lambda,\ a \leq g\}$*, and the isomorphism carries it to the orthocomplementation* $X \mapsto X^\perp$ *of* $\mathrm{DM}(\Lambda)$.

The meet determines the order through the equivalence $a \leq b$ if and only if $a \wedge b = a$, which we use in Lemma 5.4 to recover the order from the lattice operations. If every subset has a join, its meet is the join of all its lower bounds. Reversing the order gives the converse, so completeness may equally be stated in terms of meets.

Orthomodular or not, every ortholattice satisfies the De Morgan law $(\bigvee S)^\perp = \bigwedge\{a^\perp : a \in S\}$ for every subset S whose join exists, since the map $a \mapsto a^\perp$ is an order-reversing bijection that is its own inverse. It is through this law that orthocomplements of joins are computed below, in the MacNeille completion and in the coordinatization theorem of Section A.3.

For complete orthomodular lattices that are atomistic and have the covering property, irreducibility has an equivalent formulation through direct products: such a lattice is irreducible exactly when it admits no decomposition as a direct product of two lattices, each with more than one element (see Maeda and Maeda [MM70] and Holland [Hol95, §4]). We adopt the definition of irreducibility in Definition A.3 because it is the form in which irreducibility enters the coordinatization theorem (see Theorem A.16).

Let X be a cut of an ortholattice Λ. The set $X^\perp$, as a set of lower bounds, is a cut, and the map $X \mapsto X^\perp$ reverses inclusion. An element b lies in $X^\perp$ exactly when $b^\perp$ is an upper bound of X, whence $X^{\perp\perp} = (X^u)^l = X$. Moreover, the intersection $X \cap X^\perp$ is $\{0\}$, because an element b with $b \leq b^\perp$ satisfies $b = b \wedge b^\perp = 0$. Since $X \mapsto X^\perp$ is an order-reversing bijection that is its own inverse, it carries the meet $X \cap X^\perp$ to the join of $X^\perp$ and X, which is therefore $\{0\}^\perp = \Lambda$. Thus the map $X \mapsto X^\perp$ is an orthocomplementation of $\mathrm{DM}(\Lambda)$; it extends that of Λ, since it sends the cut $\{b \in \Lambda : b \leq c\}$ to the cut $\{b \in \Lambda : b \leq c^\perp\}$.

Proof of Proposition A.5. For $g \in \Gamma$, write $\Lambda_{\leq g} = \{a \in \Lambda : a \leq g\}$; we show that each $\Lambda_{\leq g}$ is a cut and that the map $g \mapsto \Lambda_{\leq g}$ is an order embedding whose image contains every cut. The assertions on orthocomplementations then follow from the De Morgan law.

By join-density, we have $g = \bigvee \Lambda_{\leq g}$; hence the upper bounds of $\Lambda_{\leq g}$ in Λ are the elements $b \in \Lambda$ with $b \geq g$. By meet-density, the common lower bounds of these elements are the elements $a \in \Lambda$ with $a \leq g$; thus the set $\Lambda_{\leq g}$ is a cut. For $g, g' \in \Gamma$, join-density also shows that $g \leq g'$ if and only if $\Lambda_{\leq g} \subseteq \Lambda_{\leq g'}$; hence the map is an order embedding. Because the order of Γ induces that of Λ, the map restricts on Λ to the canonical embedding.

If X is a cut and $g = \bigvee X$ in Γ, then the cut X is contained in $\Lambda_{\leq g}$. Conversely, an element $a \in \Lambda$ with $a \leq g$ lies below every $b \in X^u$ because each such b is an upper bound of X in Γ, whence $a \in X = (X^u)^l$. Thus $X = \Lambda_{\leq g}$, and the map is onto $\mathrm{DM}(\Lambda)$.

The formula for $g^\perp$ follows from $g = \bigvee \Lambda_{\leq g}$ and the De Morgan law. By this formula, an element $b \in \Lambda$ lies below $g^\perp$ exactly when $b \leq a^\perp$ for every $a \in \Lambda_{\leq g}$; hence $\Lambda_{\leq g^\perp} = (\Lambda_{\leq g})^\perp$. □

Historical note. Birkhoff and von Neumann rejected the lattice of closed subspaces of an infinite-dimensional Hilbert space as the logic of quantum mechanics. Letters that von Neumann wrote to Birkhoff in 1935, while the two were preparing their paper of 1936, show why: the lattice is not modular. Von Neumann, however, wanted the probabilities of quantum propositions to behave like a dimension function, a real-valued function d on the lattice that is strictly increasing and satisfies $d(a \vee b) + d(a \wedge b) = d(a) + d(b)$. Only a modular lattice carries such a function [Réd07]. In the paper itself, Birkhoff and von Neumann postulated

instead an abstract projective geometry, whose lattice of subspaces is modular [BvN36]. A year later Husimi observed that the closed subspaces of a Hilbert space satisfy the orthomodular law [Hus37], which was rediscovered several times in the 1950s and 1960s. In his thesis of 1964, Piron took completeness, orthomodularity, atomicity and the covering property as axioms for the propositions about a quantum system. He showed that the irreducible lattices with these properties are isomorphic to lattices of orthogonally closed subspaces of hermitian spaces when their dimension is at least four [Pir64].

A.2. **Hermitian spaces.** A one-dimensional subspace of a vector space is called a *ray*. By the coordinatization theorem of projective geometry, every projective geometry with four independent points is isomorphic to the geometry of the rays of a vector space over a division ring [Bae52]. The theorem has an orthocomplemented form, which represents an ortholattice with the properties listed below by the orthogonally closed subspaces of a hermitian space (Theorem A.16); the division ring that it yields need not be $\mathbb{R}$, $\mathbb{C}$ or $\mathbb{H}$.

Definition A.6. A *division ring* is a ring with $1 \neq 0$ in which every nonzero element is invertible; multiplication need not be commutative. An *involution* of a division ring Δ is a map $\alpha \mapsto \alpha^*$ such that $(\alpha+\beta)^* = \alpha^*+\beta^*$, $(\alpha\beta)^* = \beta^*\alpha^*$ and $\alpha^{**} = \alpha$ for all $\alpha, \beta \in \Delta$. Fix an involution $\alpha \mapsto \alpha^*$ of Δ, and let E be a left vector space over Δ, that is, a left Δ-module.[7] A *hermitian form* on E is a map $\langle \cdot \mid \cdot \rangle \colon E \times E \to \Delta$ such that $\langle \alpha x + \beta y \mid z \rangle = \alpha\langle x \mid z\rangle + \beta\langle y \mid z\rangle$ and $\langle y \mid x\rangle = \langle x \mid y\rangle^*$ for all $x, y, z \in E$ and $\alpha, \beta \in \Delta$; it is *anisotropic* when $\langle x \mid x\rangle \neq 0$ for every nonzero vector $x \in E$. A *hermitian space* is a left vector space over a division ring with an involution, together with an anisotropic hermitian form.

Definition A.7. In a hermitian space E, two vectors x and y are *orthogonal*, written $x \perp y$, when $\langle x \mid y\rangle = 0$. For $S \subseteq E$, let $S^\perp$ be the set $\{x \in E : \langle x \mid y\rangle = 0 \text{ for all } y \in S\}$. A subspace M is *orthogonally closed* when $M^{\perp\perp} = M$.

Definition A.8 (The lattice $\mathcal{C}(E)$). For a hermitian space E, we write $\mathcal{C}(E)$ for the set of orthogonally closed subspaces of E, ordered by inclusion.

The set $\mathcal{C}(E)$, with the map $M \mapsto M^\perp$, is a complete ortholattice, as we verify below.

Definition A.9. A hermitian space E is *orthomodular* when $M + M^\perp = E$ for every $M \in \mathcal{C}(E)$, where $M + M^\perp$ is the set of sums $v + y$ with $v \in M$ and $y \in M^\perp$.

Definition A.10. A *Hilbert space* is a hermitian space H over the real numbers $\mathbb{R}$, the complex numbers $\mathbb{C}$ or the quaternions $\mathbb{H}$, with the usual involution (the identity, complex conjugation or quaternionic conjugation), such that $\langle x \mid x\rangle > 0$ for every nonzero vector x and the space is complete in the norm $\|x\| = \langle x \mid x\rangle^{1/2}$.

Lemma A.11. *In a hermitian space E over a division ring Δ, the subspace $M + \Delta x$ belongs to $\mathcal{C}(E)$ for every $M \in \mathcal{C}(E)$ and every $x \in E$.*

It follows from the identities that define a hermitian form that $\langle x \mid \alpha y\rangle = \langle x \mid y\rangle\alpha^*$ for all $x, y \in E$ and $\alpha \in \Delta$. The bracket is linear in its first argument, as is customary in mathematics; in the Dirac notation used by physicists it is linear in the second.

Nothing in the definition of a hermitian space forces the scalars to be real, complex or quaternionic: over $\mathbb{Q}$ with the identity involution, the form $\sum_n x_n y_n$ on finitely supported rational sequences $x = (x_n)_{n\in\mathbb{N}}$ and $y = (y_n)_{n\in\mathbb{N}}$ is anisotropic, since a sum of rational squares vanishes only when every square does.

Orthogonality is symmetric, since $\langle y \mid x\rangle = \langle x \mid y\rangle^*$ and the involution fixes 0. In particular, every subset S of E satisfies $S \subseteq S^{\perp\perp}$.

[7] Physicists usually let quaternionic scalars act on the right; their quaternionic Hilbert spaces are therefore right modules. The choice of side is immaterial: a left vector space over Δ is a right vector space over the opposite division ring, and an involution is an isomorphism of Δ onto its opposite.

Every set of the form $S^\perp$ is an orthogonally closed subspace: applying the inclusion-reversing operation $^\perp$ to $S \subseteq S^{\perp\perp}$ gives $S^{\perp\perp\perp} \subseteq S^\perp$, and the reverse inclusion is the same fact for $S^\perp$.

If $(M_j)_{j\in J}$ is a family of orthogonally closed subspaces indexed by an arbitrary set J, then their intersection, which equals $\bigcap_j M_j^{\perp\perp} = \big(\bigcup_j M_j^\perp\big)^\perp$, is orthogonally closed. Hence the set $\mathcal{C}(E)$ is a complete lattice with least element $\{0\}$ and greatest element E. In this lattice the family has meet $\bigcap_j M_j$ and join $\big(\sum_j M_j\big)^{\perp\perp}$, which may be strictly larger than $\sum_j M_j$.

For $M \in \mathcal{C}(E)$, anisotropy gives $M \cap M^\perp = \{0\}$, whence $M \vee M^\perp = (M^\perp \cap M)^\perp = E$. Since the map $M \mapsto M^\perp$ also reverses inclusion and is its own inverse on $\mathcal{C}(E)$, the lattice $\mathcal{C}(E)$ is, with this map, a complete ortholattice.

In a Hilbert space H, the projection theorem supplies the decomposition that defines orthomodular spaces above: every closed subspace M satisfies $M + M^\perp = H$. For every subspace M, it also gives $M^{\perp\perp} = \overline{M}$. Thus orthogonal closure is ordinary norm closure, and $L(H) = \mathcal{C}(H)$. The space H is orthomodular, and Proposition A.13 below shows that its lattice is orthomodular as well.

Lemma A.12. *In a hermitian space E, every finite-dimensional subspace F is orthogonally closed and satisfies $F + F^\perp = E$.*

Proof. The idea is to write down the orthogonal projection onto F. Let n be the dimension of F. By anisotropy, the Gram–Schmidt process gives F an orthogonal basis $w_1, \dots, w_n$ with $\langle w_i \mid w_i\rangle \neq 0$ for $i = 1, \dots, n$. For $z \in E$, the vector $Pz = \sum_i \langle z \mid w_i\rangle \langle w_i \mid w_i\rangle^{-1} w_i$ lies in F, and the difference $z - Pz$ lies in $F^\perp$; hence $F + F^\perp = E$. If $y \in F^{\perp\perp}$, the vector $y - Py$ lies in $F^\perp \cap F^{\perp\perp}$ and is therefore orthogonal to itself; by anisotropy it is 0, whence $y \in F$. □

Proof of Lemma A.11. If $x \in M$, there is nothing to prove; we therefore assume that $x \notin M$. Since $M + \Delta x$ is contained in $(M + \Delta x)^{\perp\perp}$, it suffices to show that every vector of $(M + \Delta x)^{\perp\perp}$ differs from a multiple of x by a vector of M. We compare the functional that such a vector induces on $M^\perp$ through the form with the one induced by x. Their kernels will show that they differ by multiplication on the right by a scalar. Write $W = M^\perp$, and let ξ_x be the functional $z \mapsto \langle z \mid x\rangle$ on W. This functional is nonzero, since $x \notin W^\perp = M$, and we may choose $z_0 \in W$ with $\xi_x(z_0) = 1$.

Let y be an element of $(M + \Delta x)^{\perp\perp}$, and let ξ_y be the functional $z \mapsto \langle z \mid y\rangle$ on W. Since $W \cap x^\perp = (M + \Delta x)^\perp$, the functional ξ_y vanishes on the kernel of ξ_x; as $z - \xi_x(z) z_0$ lies in this kernel for every $z \in W$, we obtain $\xi_y(z) = \xi_x(z)\alpha$ with $\alpha = \xi_y(z_0)$. Because $\langle z \mid \alpha^* x\rangle = \langle z \mid x\rangle \alpha$, the vector $y - \alpha^* x$ is orthogonal to W and therefore lies in $W^\perp = M$, whence $y \in M + \Delta x$. □

Proposition A.13. *For a hermitian space E, the lattice $\mathcal{C}(E)$ is orthomodular if and only if E is an orthomodular space.*

Proof. Suppose first that E is orthomodular, and let $M \subseteq M'$ be elements of $\mathcal{C}(E)$. Every vector $z \in M'$ can be written as $v + y$ with $v \in M$ and $y \in M^\perp$; since $y = z - v$ lies in M', we have $M' = M + (M' \cap M^\perp) \subseteq M \vee (M' \wedge M^\perp) \subseteq M'$, which is the orthomodular law.

Conversely, suppose that $\mathcal{C}(E)$ is orthomodular, and let Δ be the division ring of scalars of E. Let M be an element of $\mathcal{C}(E)$, and let x be a vector of $E \setminus M$. By Lemma A.11, the subspace $M' = M + \Delta x$ belongs to $\mathcal{C}(E)$, and orthomodularity gives $M' = M \vee (M' \cap M^\perp)$. Since $M' \neq M$, the subspace $M' \cap M^\perp$ is nonzero. As this subspace meets M only in 0, and as M has codimension one in M', the subspace $M' \cap M^\perp$ is one-dimensional, spanned by a

vector u. Then $M \subsetneq M + \Delta u \subseteq M'$, and the codimension count gives $M' = M + \Delta u$; hence $x \in M + M^{\perp}$. □

Historical note. In 1964 Piron published a theorem asserting that every orthomodular complex inner product space is complete, but his proof contained an error [Pir64]. Amemiya and Araki gave a correct proof in 1966, and their method also covers real and quaternionic scalars [AA66]. It was then speculated that, among infinite-dimensional hermitian spaces over arbitrary division rings, orthomodularity alone characterizes the Hilbert spaces over $\mathbb{R}$, $\mathbb{C}$ and $\mathbb{H}$. Keller ended the speculation in 1980 with a "non-classical Hilbert space", an infinite-dimensional orthomodular hermitian space over a division ring other than these three [Kel80]. Keller was a student of Herbert Gross, whose papers and book on forms in infinite-dimensional vector spaces shaped the theory; the story is told in Holland's survey of 1995 [Hol95].

A.3. **Projective geometry, Solèr lattices and pre-Solèr lattices.** The lattice of closed subspaces remembers more than inclusion: its atoms are rays, their joins describe incidence, and orthocomplementation records orthogonality. The coordinatization theorem recovers a vector space and a hermitian form from this information. It gives an embedding into the lattice of orthogonally closed subspaces; completeness is needed to make the embedding onto. Solèr's condition then supplies the infinite orthonormal sequence that, together with orthomodularity of the representing space, forces real, complex or quaternionic Hilbert geometry.

We develop these steps to distinguish the Solèr lattices from the pre-Solèr lattices, which may lack completeness. In the axiomatization of Section 5, the Hilbert sort provides vectors on which lattice elements act as orthogonal projections, while the sequence sort supplies all closed spans in standard models (see Theorem 7.3). We assume familiarity with Hilbert spaces but none with projective geometry, whose incidence axioms we recall before recovering the vectors from an abstract lattice.

A.3.1. *Rays and coordinates.* In a complex Hilbert space H, the atoms of $L(H)$ are the rays $\mathbb{C}x$ spanned by nonzero vectors x. Two distinct rays have as their join the plane they span, and the rays in the plane are exactly the atoms below the join (see Figure 10). Thus the lattice records which rays lie in which planes without naming vectors or scalars.

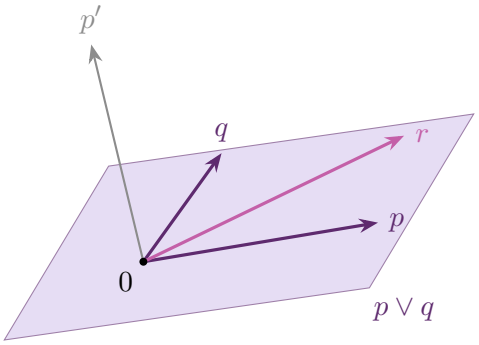


FIGURE 10. The join $p \vee q$ of two rays p and q in the lattice $L(H)$ of closed subspaces of a complex Hilbert space H, which is the plane they span and contains the ray r but not the ray p'.

To recover vectors from an abstract lattice, we first describe incidence without assuming a vector space. Coordinatization will identify this abstract incidence structure with the geometry of rays of a vector space.

Definition A.14. A *projective geometry* is a set of *points* together with a family of subsets, called *lines*, such that:

(PG1) any two distinct points lie on exactly one line;
(PG2) every line contains at least three points;
(PG3) (*the Veblen–Young axiom*) if a, b, a', b' are distinct points and the line through a and b meets the line through a' and b', then the line through a and a' meets the line through b and b'.

A set of points is a *projective subspace* when it contains the whole line through any two of its points. The projective subspace *generated* by a set of points is the intersection of all projective subspaces containing it. A *projective plane* is the projective subspace generated by three points not on one line. Axiom (PG3) says that two lines lying in a common projective plane always meet (see Figure 11); a projective geometry has no parallel lines. The word *projective* comes from perspective drawing, in which parallel rails, projected onto the canvas, meet at the horizon.

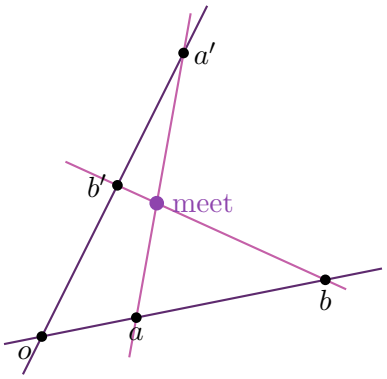


FIGURE 11. The Veblen–Young axiom: the lines ab and $a'b'$ meet at o, and then the lines aa' and bb' meet as well.

The basic example comes from linear algebra. Let V be a left vector space of dimension at least two over a division ring Δ. Take as points the rays of V, and as lines the sets of rays contained in a two-dimensional subspace. Axiom (PG1) holds because two distinct rays span a unique two-dimensional subspace. Axiom (PG2) holds because the two-dimensional subspace spanned by linearly independent vectors x and y contains the three distinct rays Δx, Δy and $\Delta(x + y)$. For (PG3), suppose that the two-dimensional subspaces spanned by a, b and by a', b' have a ray in common. The rays a, b, a', b' therefore lie in a subspace of dimension at most three. By counting dimensions, any two two-dimensional subspaces of this subspace have a ray in common, in particular those spanned by a, a' and by b, b'.

Points are *independent* when none of them lies in the projective subspace generated by the others. In the basic example, rays are independent exactly when nonzero vectors spanning them, one in each ray, are linearly independent. Two projective geometries are *isomorphic* when there is a bijection between their sets of points under which the lines of the one correspond exactly to the lines of the other. For vector spaces E and E' over division rings Δ and Δ', respectively, a *semilinear isomorphism* from E onto E' is an additive bijection $f\colon E \to E'$

for which there is an isomorphism $\sigma\colon \Delta \to \Delta'$ with $f(\alpha x) = \sigma(\alpha) f(x)$ for all $\alpha \in \Delta$ and $x \in E$. The following theorem says that the basic example is, up to isomorphism, the only one as soon as there are four independent points.

Theorem A.15 (Coordinatization of projective geometries; see Baer [Bae52]). *Every projective geometry that contains four independent points is isomorphic to the geometry of rays of a left vector space E over a division ring Δ. The pair (Δ, E) is unique up to a semilinear isomorphism.*

The division ring Δ is thus recovered from the geometry alone, whatever it may be: every division ring occurs. The hypothesis of four independent points, on the other hand, is indispensable: a projective plane regarded as a projective geometry in its own right need not come from a division ring. The planes that do are exactly those in which Desargues's theorem holds (if two triangles are in perspective from a point, the three intersection points of corresponding sides lie on one line). Planes in which Desargues's theorem fails, the *non-Desarguesian* planes, exist; in a geometry with four independent points, by contrast, the theorem follows from the axioms.

A.3.2. *Recovering a hermitian space from the lattice.* The following properties of an ortholattice Λ enter the representation theorems:

(C) the lattice is complete;
(OM) it is orthomodular;
(A) it is atomistic;
(CP) it has the covering property;
(I) it is irreducible;
(D) it contains four mutually orthogonal atoms.

Theorem A.16 (Coordinatization; Baer [Bae52], Maeda and Maeda [MM70, §34], Holland [Hol95, §3]). *Let Λ be an ortholattice that satisfies (OM), (A), (CP), (I) and (D), but not necessarily (C). Then there are a division ring Δ with an involution, a hermitian space E over Δ, and a bijection $p \mapsto \hat{p}$ from the atoms of Λ onto the rays of E, such that two atoms p and q are orthogonal exactly when $\hat{p} \perp \hat{q}$. The map*

$$\gamma\colon \Lambda \to \mathcal{C}(E), \qquad \gamma(a) = \{0\} \cup \bigcup \{\hat{p} : p \text{ an atom with } p \le a\},$$

is an embedding of ortholattices whose image contains every finite-dimensional subspace of E and the orthogonal complement of each. The triple formed by Δ, E and the form is unique up to a semilinear isomorphism followed by a rescaling $\langle x \mid y\rangle' = \langle x \mid y\rangle\zeta$, where $\zeta^* = \zeta \neq 0$, with the new involution $\alpha \mapsto \zeta^{-1}\alpha^*\zeta$.

To apply Theorem A.15 to an orthomodular lattice that is atomistic and has the covering property, we take the atoms of the lattice as the points of a geometry. The lines of this geometry are the sets $\{r : r \text{ an atom with } r \le p \vee q\}$ for distinct atoms p and q; in $L(H)$, these sets consist of the rays in the plane spanned by p and q. The covering property gives this geometry (PG1) and, with orthomodularity, (PG3); irreducibility in the sense of Definition A.3 is exactly (PG2) (see Maeda and Maeda [MM70, §§15–16] and Holland [Hol95, §4]).

Orthocomplementation adds a relation between points: by the definition of orthogonality, atoms p and q are orthogonal when $p \le q^\perp$. In $L(H)$ this is orthogonality of rays, and it determines the inner product up to the ambiguity described at the end of Theorem A.16. Condition (D) supplies the four independent points of Theorem A.15, since mutually orthogonal atoms are independent. The coordinatization theorem, Theorem A.16, goes back, in finite dimensions, to Birkhoff and von Neumann [BvN36]. The theorem carries the atoms

onto all rays of a hermitian space and preserves their orthogonality, while embedding the lattice into the orthogonally closed subspaces of the space. Condition (C) is absent from its hypotheses, so the embedding applies to the incomplete lattices of Section 6; surjectivity requires a separate argument, given in the proof of Theorem A.18.

Proof of the statements of Theorem A.16 about γ. By Theorem A.15, incidence among the atoms recovers a vector space E whose rays correspond to all the atoms. Orthocomplementation induces a *polarity* in the sense of Holland [Hol95, Definition 3.2], sending a projective subspace to the points orthogonal to all its points. It reverses inclusion and sends the whole geometry to the empty set. The polar of a point p consists of the atoms below $p^\perp$. It is a hyperplane, or maximal proper projective subspace, whose polar is $\{p\}$ [Hol95, §4].

The projective-dual construction for Baer's representation [Bae52, Chapter IV] gives a form B inducing the polarity, linear in its first variable and g-semilinear in its second for an anti-automorphism g of Δ (see Holland [Hol95, proof of Theorem 3.6]). Thus $B(x,y)=0$ for nonzero vectors $x,y\in E$ exactly when the corresponding atoms are orthogonal. The form is reflexive, meaning that $B(x,y)=0$ if and only if $B(y,x)=0$, and it is anisotropic because $p\wedge p^\perp=0$ for every atom p. To obtain a hermitian form, choose $e\neq 0$, let $d=B(e,e)$, and define

$$B_0(x,y)=B(x,y)d^{-1},\qquad \tau(\alpha)=d\,g(\alpha)d^{-1}.$$

The forms B_0 and B have the same zero set, so B_0 induces the same polarity. Moreover, B_0 is hermitian for the involution τ, in every characteristic.[8] We verify that the resulting map γ represents the whole lattice and preserves its operations. If two atoms lie below an element a of Λ, every atom on the line they determine lies below a; hence the set $\gamma(a)$ is a subspace. For an atom r the following are equivalent: $r\le a^\perp$; $a\le r^\perp$; $p\le r^\perp$ for every atom $p\le a$ (because (A) makes a the join of these atoms); $\hat r\perp\hat p$ for every atom $p\le a$. Thus the map γ carries $a^\perp$ to $\gamma(a)^\perp$, and applying this twice gives $\gamma(a)=\gamma(a)^{\perp\perp}$. By (A), elements $a,b\in\Lambda$ satisfy $a\le b$ if and only if $\gamma(a)\subseteq\gamma(b)$, which makes γ injective. Meets are preserved because an atom lies below $a\wedge b$ exactly when it lies below a and b, and joins are preserved by the De Morgan law. Finally, let F be a finite-dimensional subspace of E with basis $x_1,\dots,x_k$, and let p_i be the atom with $\hat p_i=\Delta x_i$, for $i=1,\dots,k$. The join $p_1\vee\cdots\vee p_k$ exists, as every finite join does in a lattice. Because γ preserves joins, the image of this join is $(\Delta x_1+\cdots+\Delta x_k)^{\perp\perp}=F^{\perp\perp}$, which is F by Lemma A.12. The orthogonal complement $F^\perp$ is therefore $\gamma\big((p_1\vee\cdots\vee p_k)^\perp\big)$. □

Piron [Pir64] proved this representation for complete lattices; completeness is not needed for the embedding, and in the proof of Theorem A.18 below it serves only to make γ surjective.

A.3.3. *Harmonic conjugates and Solèr's condition.* Solèr proved in 1995 that an orthomodular space containing an infinite orthonormal sequence is a Hilbert space over $\mathbb{R}$, $\mathbb{C}$ or $\mathbb{H}$ [Sol95], and Holland restated the existence of such a sequence as a condition on the lattice, through harmonic conjugates, a construction of classical projective geometry that uses joins and meets

[8] For $y\neq 0$, the left-linear functionals $L_y(x)=B_0(x,y)$ and $H_y(x)=\tau^{-1}(B_0(y,x))$ have the same kernel, so $H_y=L_y\lambda_y$ for a right scalar λ_y. For independent y,z, the functionals L_y and L_z are independent over the scalars acting on the right, since a relation between them would, by the nondegeneracy of B_0, give a relation between y and z; comparing $H_{y+z}=H_y+H_z=L_y\lambda_y+L_z\lambda_z$ with $H_{y+z}=L_{y+z}\lambda_{y+z}=L_y\lambda_{y+z}+L_z\lambda_{y+z}$ therefore gives $\lambda_y=\lambda_z=\lambda_{y+z}$. When y and z are dependent, an auxiliary vector independent of both gives the same equality, so λ_y is constant for $y\neq 0$. Since $B_0(e,e)=1$, this constant is 1. Hence $B_0(y,x)=\tau(B_0(x,y))$; applying this twice and using $B_0(e,\alpha e)=\tau(\alpha)$ gives $\tau^2=\mathrm{id}$.

alone [Hol95]. This condition, together with completeness and the conditions of the coordinatization theorem, singles out the lattices of closed subspaces of infinite-dimensional Hilbert spaces (Theorem A.18).

Let E be a left vector space of dimension at least three over a division ring Δ. For distinct rays p and q of E, regarded as projective points, write $p + q$ for their projective line. Fix a point c on this line other than p and q. Choose a point r not on $p + q$. All intersections below lie in the ray geometry of the three-dimensional subspace spanned by p, q, r, where two distinct projective lines meet in one point; we denote the point by their intersection. Choose a point r' on $r + p$ other than r and p; let $c' = (c + r') \cap (q + r)$ and $c'' = (p + c') \cap (q + r')$. The point $h = (r + c'') \cap (p + q)$ is the *harmonic conjugate* of c with respect to p and q. The point h does not depend on the choices of r and r', and, for vectors u and v spanning p and q, respectively,

$$\text{the harmonic conjugate of } \Delta(u + v) \text{ with respect to } \Delta u \text{ and } \Delta v \text{ is } \Delta(u - v). \tag{1}$$

The bijection of Theorem A.16 between atoms and rays transfers the harmonic conjugate to the atoms of an ortholattice that satisfies the hypotheses of that theorem.

Solèr's condition, in the form we use, asks for an infinite orthogonal sequence of atoms along which the orthocomplement and the harmonic conjugate agree at every step:

(S) There are mutually orthogonal atoms p_n, for $n \in \mathbb{N}$, and atoms $c_n \leq p_n \vee p_{n+1}$ with $c_n \neq p_n, p_{n+1}$, such that, for every $n \in \mathbb{N}$, the harmonic conjugate of c_n with respect to p_n and p_{n+1} is $c_n^\perp \wedge (p_n \vee p_{n+1})$.

Definition A.17. A *Solèr lattice* is an ortholattice that satisfies (C), (OM), (A), (CP), (I), (D) and (S). A *pre-Solèr lattice* is an ortholattice that satisfies (OM), (A), (CP), (I), (D) and (S), but not necessarily (C).

Theorem A.18 (Piron, Amemiya and Araki, Solèr; in this form Holland [Hol95, Theorem 4.1]**).** *An ortholattice is a Solèr lattice if and only if it is isomorphic to $L(H)$ for some infinite-dimensional Hilbert space H over $\mathbb{R}$, $\mathbb{C}$ or $\mathbb{H}$.*

Coordinatization leaves the division ring unrestricted. For example, the lattice of all subspaces of $\mathbb{Q}^4$, orthocomplemented by the form $\sum_{i<4} x_i y_i$, satisfies the hypotheses of Theorem A.16 with $\Delta = \mathbb{Q}$. To reach Hilbert spaces, we need a further condition on the lattice that supplies an infinite orthonormal sequence in its representing space.

Solèr's condition does this through harmonic conjugates. For two orthogonal rays, the harmonic conjugate of a third ray in their plane agrees with its orthocomplement in the plane exactly when the third ray is spanned by the sum of two vectors of equal "length", one from each orthogonal ray. Applying this requirement along an infinite orthogonal sequence permits a single rescaling of the form that makes the sequence orthonormal (see Lemma A.19). Completeness of the lattice will supply orthomodularity of the whole representing space, the other hypothesis needed for Solèr's theorem.

The harmonic conjugate is defined above, and Figure 12 shows its construction. In the vector space E, the construction has a simple coordinate description, which shows that the harmonic conjugate does not depend on the choices of the auxiliary points r and r'. Let u and v be vectors spanning p and q, respectively, scaled so that $c = \Delta(u + v)$; this is possible because c lies on $p+q$ and differs from p and q. Likewise, let w be a vector spanning r, scaled so that $r' = \Delta(u + w)$, which is possible because r' lies on $r + p$ and differs from r and p. A direct computation gives $c' = \Delta(v - w)$, $c'' = \Delta(u - v + w)$ and $h = \Delta(u - v)$. All coefficients are ± 1, so the computation is valid over every division ring, which proves (1).

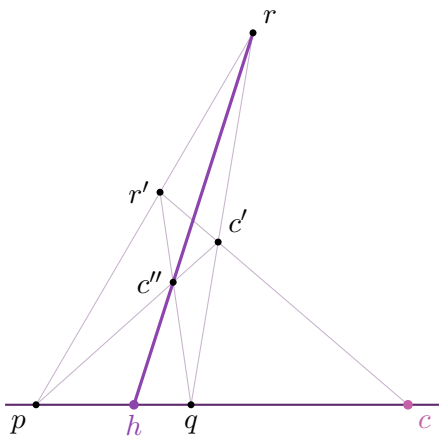


FIGURE 12. The construction of the harmonic conjugate h of a point c of the line through p and q, with respect to p and q, from the auxiliary points r, r', c' and c''.

Now let Λ be an ortholattice satisfying the hypotheses of Theorem A.16, and apply the construction in its coordinatizing space E. The bijection between atoms and rays transfers the harmonic conjugate to the atoms, with $\vee$ in place of $+$ and $\wedge$ in place of $\cap$. The coordinate calculation proves independence of the auxiliary choices, so the relation "h is the harmonic conjugate of c with respect to p and q" is expressed by a first-order formula in the language of lattices.

Suppose now that the vector space carries an anisotropic hermitian form, as the space E of Theorem A.16 does, and that $p \perp q$; then the vectors u and v are orthogonal. The orthocomplement of c within the plane $p \vee q$, the ray $c^\perp \wedge (p \vee q)$, consists of the vectors $\alpha u + \beta v$ with $\langle \alpha u + \beta v \mid u + v \rangle = \alpha \langle u \mid u \rangle + \beta \langle v \mid v \rangle = 0$. It coincides with the harmonic conjugate $\Delta(u - v)$ exactly when $\langle u \mid u \rangle = \langle v \mid v \rangle$, that is, when u and v have the same "length" (see Figure 13).

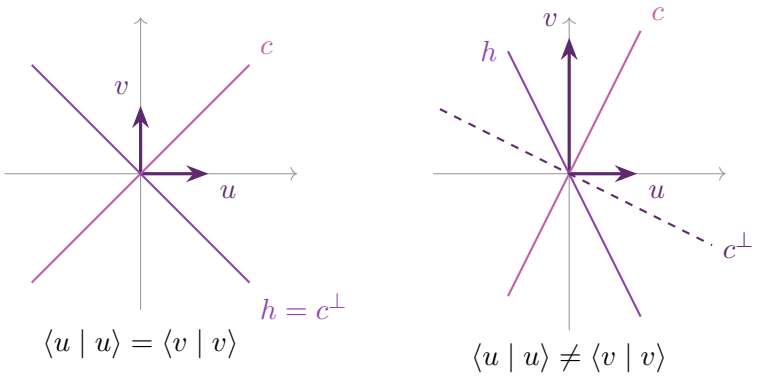


FIGURE 13. The ray $c = \Delta(u + v)$ spanned by the sum of orthogonal vectors u and v, its harmonic conjugate $h = \Delta(u - v)$ with respect to Δu and Δv, and its orthocomplement $c^\perp$ in the plane of u and v, when u and v have equal length (left) and when they do not (right).

This formulation of Solèr's condition through harmonic conjugates is due to Holland [Hol95, Theorems 3.7 and 4.1]. Morash's angle-bisecting axiom [Mor73], which asks for bisectors as in the left half of Figure 13, may be used instead (see Holland [Hol95, §4]). The next lemma verifies the role of (S): agreement of the harmonic conjugate with the orthocomplement provides a common normalization of the infinite orthogonal sequence.

Lemma A.19. *Let Λ be an ortholattice that satisfies (OM), (A), (CP), (I), (D) and (S), and let E be a hermitian space over a division ring Δ and $p \mapsto \hat{p}$ a bijection from the atoms of Λ onto the rays of E that Theorem A.16 provides for Λ. After a rescaling of the form, the space E contains an infinite orthonormal sequence.*

Proof. Let p_n and c_n, for $n \in \mathbb{N}$, be atoms as in (S). Choose e_0 with $\hat{p}_0 = \Delta e_0$, and inductively choose e_{n+1} with $\hat{p}_{n+1} = \Delta e_{n+1}$ and $\hat{c}_n = \Delta(e_n + e_{n+1})$. The choice of e_{n+1} is possible because, by Theorem A.16, the ray $\hat{c}_n$ lies in the plane $\hat{p}_n + \hat{p}_{n+1}$ and differs from $\hat{p}_n$ and $\hat{p}_{n+1}$. By (1) and the computation of the orthocomplement of c in the plane $p \vee q$ above, condition (S) gives $\langle e_n \mid e_n\rangle = \langle e_{n+1} \mid e_{n+1}\rangle$ for every $n \in \mathbb{N}$; call this common value ζ, which satisfies $\zeta^* = \zeta \neq 0$. The form $\langle x \mid y\rangle' = \langle x \mid y\rangle\zeta^{-1}$ is hermitian for the involution $\alpha \mapsto \zeta\alpha^*\zeta^{-1}$, has the same orthogonality relation, and satisfies $\langle e_n \mid e_n\rangle' = 1$ for every n; as the atoms p_n are mutually orthogonal, the sequence $(e_n)_{n\in\mathbb{N}}$ is orthonormal for the new form. □

Sketch of the proof of Theorem A.18. Let Λ be a Solèr lattice, with γ and E as in Theorem A.16. Given $M \in \mathcal{C}(E)$, completeness provides the join a of the atoms p with $\hat{p} \subseteq M$, and this join satisfies $\gamma(a)^\perp = \gamma(a^\perp) = M^\perp$, whence $\gamma(a) = M$. Thus the embedding γ is an isomorphism onto $\mathcal{C}(E)$; the lattice $\mathcal{C}(E)$ is therefore orthomodular, and by Proposition A.13 so is the space E. By Lemma A.19, the space E contains, after a rescaling of the form that changes neither $\mathcal{C}(E)$ nor its orthocomplementation, an infinite orthonormal sequence. Solèr's theorem [Sol95] (see also Prestel [Pre95]) asserts that an orthomodular space containing an infinite orthonormal sequence is a Hilbert space over $\mathbb{R}$, $\mathbb{C}$ or $\mathbb{H}$. The space E, with its form so rescaled, is therefore such a Hilbert space. Conversely, the lattice $L(H)$ is complete and satisfies the other conditions. For (S), take an orthonormal sequence $(e_n)_{n\in\mathbb{N}}$ in H, and let p_n be the ray of e_n and c_n the ray of $e_n + e_{n+1}$. □

Condition (S) cannot be omitted from Theorem A.18. Keller [Kel80] constructed infinite-dimensional orthomodular hermitian spaces over division rings other than $\mathbb{R}$, $\mathbb{C}$ and $\mathbb{H}$. For such a space E the lattice $\mathcal{C}(E)$ satisfies (C), (OM), (A), (CP), (I) and (D) (see Proposition A.13 and Lemmas A.11 and A.12). Nevertheless, by the uniqueness clause of Theorem A.16, the lattice $\mathcal{C}(E)$ is isomorphic to no lattice $L(H)$ of a Hilbert space H over $\mathbb{R}$, $\mathbb{C}$ or $\mathbb{H}$.

⧖ **Historical note.** Von Staudt's *Geometrie der Lage* of 1847 was the first treatment of projective geometry free of any metric. His later calculus of "throws" outlined the algebraization of an axiomatic geometry. Hilbert showed in his *Grundlagen der Geometrie* of 1899 that Desargues's theorem follows from the axioms of space but cannot be proved in the plane without the axioms of congruence. On this theorem he founded an algebra of segments that does not use the axioms of congruence [Hil99].

Veblen and Young gave axioms for projective geometry in 1908 [VY08]. In their book of 1910 they introduced coordinates through an algebra of points that descends from von Staudt's algebra of throws [VY10]. Solèr, a student of Herbert Gross, had just begun work on Gross's orthomodular problem when Gross died suddenly in 1989. She completed the work under the informal guidance of Alexander Prestel of Konstanz, a friend of Gross's, and

her theorem appeared in 1995 [Sol95]. In a survey of the same year, Holland restated her condition through harmonic conjugates, constructed as in the book of Veblen and Young, in the form used in this subsection [Hol95].

A.4. **What a pre-Solèr lattice is missing.** The preceding subsection used completeness to pass from a pre-Solèr lattice to the full lattice of closed subspaces of a Hilbert space. We now identify its MacNeille completion with the ortholattice of orthogonally closed subspaces of the hermitian space obtained by coordinatization, and show that this completion is a Solèr lattice exactly when it is orthomodular. A condition on sets of atoms characterizes orthomodularity of the completion without referring to the representing space. Neither this condition nor completeness is expressible among pre-Solèr lattices by a sentence of two-valued $L_{\omega_1\omega}$ about lattice elements (see Proposition A.24).

A pre-Solèr lattice need not be complete, as Example A.21 below shows; the following criterion and that example concern the MacNeille completion of such a lattice.

Theorem A.20. *Let Λ be a pre-Solèr lattice, and let E be a hermitian space that Theorem A.16 provides for Λ. Then the following are equivalent:*

(1) the MacNeille completion of Λ is orthomodular;
(2) the space E is orthomodular;
(3) the lattice Λ is isomorphic to a join-dense and meet-dense sub-ortholattice of $L(H)$ for some infinite-dimensional Hilbert space H over $\mathbb{R}$, $\mathbb{C}$ or $\mathbb{H}$.

If these conditions hold, then the space E, with its form suitably rescaled, is a Hilbert space as in (3). The lattice Λ is isomorphic to $L(H)$ for some infinite-dimensional Hilbert space H over $\mathbb{R}$, $\mathbb{C}$ or $\mathbb{H}$ if and only if Λ is complete.

Example A.21. Let E_0 be the space of finitely supported complex sequences $x = (x_n)_{n\in\mathbb{N}}$, with the inner product $\langle x \mid y\rangle = \sum_n x_n\overline{y_n}$ of the space ℓ^2 of square-summable sequences, and let Λ_0 be the set of finite-dimensional subspaces of E_0 and their orthogonal complements in E_0. Then Λ_0 is a pre-Solèr lattice whose MacNeille completion is isomorphic to $\mathcal{C}(E_0)$ and is not orthomodular.

Orthomodularity of the completion of a pre-Solèr lattice can be stated through the atoms of the lattice alone, without reference to the completion or to a coordinatizing space. For a set $\mathcal{A}$ of atoms, let $\mathcal{A}^\perp$ be the set of atoms orthogonal to every element of $\mathcal{A}$, and call $\mathcal{A}$ *orthoclosed* when $\mathcal{A} = \mathcal{A}^{\perp\perp}$. Consider the following condition, named (MOM) for MacNeille orthomodularity:

(MOM) for every orthoclosed set $\mathcal{A}$ of atoms and every atom $p \notin \mathcal{A}$, some atom $q \in \mathcal{A}^\perp$ lies in $(\mathcal{A} \cup \{p\})^{\perp\perp}$.

Theorem A.22. *The MacNeille completion of a pre-Solèr lattice is orthomodular if and only if the lattice satisfies (MOM).*

Throughout this subsection we fix a pre-Solèr lattice Λ, together with a hermitian space E and an embedding $\gamma\colon \Lambda \to \mathcal{C}(E)$ that Theorem A.16 provides for Λ. We identify Λ with its image $\gamma(\Lambda) \subseteq \mathcal{C}(E)$.

Theorem A.23. *Let Λ be a pre-Solèr lattice, let E be a hermitian space and $\gamma\colon \Lambda \to \mathcal{C}(E)$ an embedding that Theorem A.16 provides for Λ, and identify Λ with its image $\gamma(\Lambda)$. Then Λ is join-dense and meet-dense in $\mathcal{C}(E)$, and $\mathcal{C}(E)$, with $M \mapsto M^\perp$, is isomorphic as an ortholattice to the MacNeille completion of Λ.*

Proof. Every $M \in \mathcal{C}(E)$ is the join in $\mathcal{C}(E)$ of the rays it contains, which belong to Λ by Theorem A.16. If $M^\perp = \{0\}$, then the subspace M is E, which belongs to Λ. Otherwise M

equals $M^{\perp\perp} = \bigcap\{y^\perp : 0 \neq y \in M^\perp\}$, a meet of hyperplanes, each of which belongs to Λ. Since γ is an embedding of ortholattices, the second assertion now follows from Proposition A.5. □

The completion supplies the orthogonally closed subspaces absent from Λ. Any such subspace M must have both M and $M^\perp$ infinite-dimensional, since Λ already contains every finite-dimensional subspace and its orthogonal complement (see Theorem A.16). Theorem A.20 gives the criterion for the completion to be a Solèr lattice, and Example A.21 exhibits a pre-Solèr lattice for which the criterion fails.

Proof of Theorem A.20. Conditions (1) and (2) are equivalent by Theorem A.23 and Proposition A.13. If (2) holds, then Lemma A.19 and Solèr's theorem (see the proof of Theorem A.18) show that the space E, with its form rescaled, is a Hilbert space. Rescaling changes neither $\mathcal{C}(E)$ nor its orthocomplementation, and condition (3) follows from Theorem A.23. If (3) holds, then by Proposition A.5 the orthomodular lattice $L(H)$ is isomorphic as an ortholattice to the MacNeille completion of Λ; this is (1). For the last assertion, a complete pre-Solèr lattice is a Solèr lattice and by Theorem A.18 is therefore isomorphic to $L(H)$ for an infinite-dimensional Hilbert space H over $\mathbb{R}$, $\mathbb{C}$ or $\mathbb{H}$. Conversely, every lattice $L(H)$ is complete. □

In Example A.21, the finite-dimensional subspaces of an incomplete complex inner product space and their orthogonal complements form a pre-Solèr lattice whose MacNeille completion is not orthomodular.

Verification of Example A.21. By Lemma A.12, the finite-dimensional elements of Λ_0 are orthogonally closed; the others are orthogonal complements, hence orthogonally closed as well, and Λ_0 is closed under $^\perp$. Moreover, for finite-dimensional F and G the joins $F \vee G = F + G$, $F^\perp \vee G^\perp = (F \cap G)^\perp$ and $F \vee G^\perp = (F^\perp \cap G)^\perp$ belong to Λ_0, whence the De Morgan law places the meets in Λ_0 as well. Thus Λ_0 is a sub-ortholattice of $\mathcal{C}(E_0)$. The lattice Λ_0 is orthomodular, because every $M \in \Lambda_0$ satisfies $M + M^\perp = E_0$ by Lemma A.12, and the argument of Proposition A.13 applies. As a sub-ortholattice of $\mathcal{C}(E_0)$ that contains every ray of E_0, the lattice Λ_0 inherits (A), (CP) and (I) from $\mathcal{C}(E_0)$, in which (CP) holds by Lemma A.11. Conditions (D) and (S) hold with $p_n = \mathbb{C}e_n$ and $c_n = \mathbb{C}(e_n + e_{n+1})$, where $(e_n)_{n\in\mathbb{N}}$ is the sequence of standard unit vectors. Consequently Λ_0 is a pre-Solèr lattice, and the space E_0, with $\hat{p} = p$ for every atom p of Λ_0, is a hermitian space that Theorem A.16 provides for Λ_0. The corresponding embedding γ is then the inclusion of Λ_0 in $\mathcal{C}(E_0)$. The MacNeille completion of Λ_0, which is $\mathcal{C}(E_0)$ by Theorem A.23, is not orthomodular. Indeed, by the Amemiya–Araki theorem [AA66], a complex inner product space whose lattice of orthogonally closed subspaces is orthomodular is complete, which E_0 is not. In sum, the lattice Λ_0 is orthomodular while its MacNeille completion is not; Harding [Har91] studies such lattices. □

Proof of Theorem A.22. Let $\hat{p}$ denote, for each atom p, the ray of E that corresponds to p under Theorem A.16. The map $\mathcal{A} \mapsto \{0\} \cup \bigcup\{\hat{p} : p \in \mathcal{A}\}$ is a bijection from the orthoclosed sets of atoms onto $\mathcal{C}(E)$, carrying $\mathcal{A}^\perp$ to the orthogonal complement of the image of $\mathcal{A}$. If M is the image of $\mathcal{A}$ and $\hat{p} = \Delta x$, where Δ is the division ring of scalars of E, then by Lemma A.11 the set $(\mathcal{A} \cup \{p\})^{\perp\perp}$ corresponds to $M + \Delta x$. Condition (MOM) therefore says that $(M + \Delta x) \cap M^\perp \neq \{0\}$ whenever $M \in \mathcal{C}(E)$ and $x \notin M$. If $y = v + \alpha x$ is nonzero, with $v \in M$ and $y \in M^\perp$, then the scalar α is nonzero because $y \notin M$, whence the vector $x = \alpha^{-1}(y - v)$ lies in $M + M^\perp$. Conversely, if $x = v + y$ with $v \in M$ and $y \in M^\perp$, then the

vector $y = x - v$ is a nonzero element of $(M + \Delta x) \cap M^\perp$. Thus (MOM) is equivalent to the orthomodularity of E, and the theorem follows from Theorem A.20. □

By Lemma A.12, the requirement in (MOM) holds automatically unless the subspaces of E that correspond to $\mathcal{A}$ and to $\mathcal{A}^\perp$ in the proof of Theorem A.22 are both infinite-dimensional.

Historical note. MacNeille introduced the completion by cuts in 1937 [Mac37], generalizing the cuts by which Dedekind had constructed the real numbers from the rationals in 1872 [Ded72]. Banaschewski and Schmidt characterized it in 1956 by join-density and meet-density [Ban56, Sch56]. In 1969 Adams solved negatively Problem 36 of Birkhoff's *Lattice Theory*, in its third edition of 1967, by giving an orthocomplemented modular lattice whose completion by cuts is not orthomodular [Ada69]. His lattice consists of the finite-dimensional subspaces of an incomplete complex inner product space and their orthogonal complements, as in Example A.21, and his argument also rests on the Amemiya–Araki theorem. Harding gave in 1991, without recourse to the Amemiya–Araki theorem, a construction of orthomodular lattices whose completions by cuts are not orthomodular. Following a method that Kalmbach introduced in 1977, he glued together the Boolean algebras generated by the chains of a lattice [Har91]. Harding noted that whether every orthomodular lattice can be embedded in a complete one was still open.

A.5. **Nonexpressibility over pre-Solèr lattices.** Completeness is the one condition of a Solèr lattice that a pre-Solèr lattice may fail, and among pre-Solèr lattices no sentence of $L_{\omega_1\omega}$ about the lattice alone expresses it. In classical two-valued $L_{\omega_1\omega}$, the same holds for condition (MOM), which by Theorem A.22 characterizes the pre-Solèr lattices whose MacNeille completion is orthomodular; (MOM) quantifies over sets of atoms, whereas the logic quantifies only over individual elements of the lattice. For both conditions the obstruction is countability, as the proof of the following proposition shows.

Proposition A.24. *Neither (C) nor (MOM) is equivalent, among pre-Solèr lattices, to a sentence of two-valued $L_{\omega_1\omega}$ in the vocabulary of ortholattices expanded by constants for atoms p_n and c_n, for $n \in \mathbb{N}$, that witness (S).*

Proof. Suppose that a sentence φ of $L_{\omega_1\omega}$ were equivalent to (MOM) among pre-Solèr lattices. By Theorem A.22, the same sentence would then be equivalent to orthomodularity of their MacNeille completions. The axioms of ortholattices and the conditions (OM), (A), (CP), (I), (D) and (S) are expressed by countably many sentences of $L_{\omega_1\omega}$ in this vocabulary. These sentences hold in $L(\ell^2)$, with $p_n = \mathbb{C}e_n$ and $c_n = \mathbb{C}(e_n + e_{n+1})$ for the standard unit vectors e_n, $n \in \mathbb{N}$. Indeed, by (1), the harmonic conjugate of c_n with respect to p_n and p_{n+1} is $\mathbb{C}(e_n - e_{n+1})$, which is $c_n^\perp \wedge (p_n \vee p_{n+1})$. The sentence φ also holds in this lattice, since the lattice is complete and orthomodular and is therefore its own MacNeille completion. By the downward Löwenheim–Skolem theorem for a countable fragment, these sentences hold in a countable structure Λ'. Thus Λ' is a pre-Solèr lattice whose MacNeille completion is orthomodular. By Theorem A.20, a hermitian space E provided for Λ' by Theorem A.16 is, after rescaling its form, an infinite-dimensional Hilbert space. By Theorem A.16, the atoms of Λ' correspond to all the rays of E. There are uncountably many such rays, contradicting the countability of Λ'.

If instead φ expressed completeness among pre-Solèr lattices, the same argument would give a countable complete pre-Solèr lattice Λ'. By Theorem A.18, it would be isomorphic to $L(H)$ for an infinite-dimensional real, complex or quaternionic Hilbert space H, again contradicting countability. □

The argument for completeness needs only the downward Löwenheim–Skolem theorem and Theorem A.18, and it applies equally in real-valued $L_{\omega_1\omega}$, where equality of lattice elements is represented by a two-valued predicate. Take a countable fragment containing sentences that express the ortholattice axioms and the conditions defining pre-Solèr lattices, together with a sentence proposed to express completeness. The expansion of $L(\ell^2)$ used in the proof has a countable substructure in which every sentence of this fragment keeps its value. The quotient of this substructure by the equality predicate is a countable pre-Solèr lattice, so by Theorem A.18 it cannot be complete. This sentence takes the same value in this quotient as in the complete lattice $L(\ell^2)$, and so cannot express completeness. Every countable set of sentences about the lattice alone thus keeps its values in some countable structure, whereas the representation of a Solèr lattice requires all the rays of an infinite-dimensional Hilbert space.

Appendix B. The space of Henkin types and the model existence game

In a position of the game of Section 3, finitely many commitments and records describe a structure that is still being built. In this appendix we attach to each position the set of Henkin types compatible with it, in the space of Henkin types of Section 2.2. Josh's truthful answers and Olivia's terminal obstructions then acquire a topological meaning, and the positions and rules of the game can be compared with Keisler's forcing conditions and $\forall\exists$ formulas. We close by setting the Baire category argument behind Olivia's bookkeeping, which takes place in the space of plays, beside the argument by which Aubone, Dueñez and Iovino omit types in the space of Henkin types [ADI26]; the two arguments take place in different spaces, and their relation is an analogy, not an equivalence.

The space of Henkin types. We retain the setting of Section 3: the vocabulary $\mathcal{L}$ is countable, and T is a nonempty countable theory of closed conditions $\psi \le r$ or $\psi \ge r$ for a sentence ψ and a rational number r. The fragment $\mathcal{F}$ generated by the sentences ψ of the axioms of T is therefore countable (see [ADI26, Section 6]). We take as the new constants C of Section 2.2 the witnesses $\mathcal{W}$ of the game, countably many of each sort. The fragment $\mathcal{F}_C$ that $\mathcal{F}$ generates in $\mathcal{L}_{\mathcal{W}}$ then has a countable set $S_C^{\mathcal{F}}$ of sentences, and this set contains every sentence that occurs in a play of the game $\mathrm{G}(T)$ of Definition 3.2, since such a sentence is an atomic sentence of $\mathcal{L}_{\mathcal{W}}$ or arises from a subformula of an axiom by the substitution of closed terms. We write $\prod S_C^{\mathcal{F}}$ for the product of the closed intervals $[-r_\varphi, r_\varphi]$, for $\varphi \in S_C^{\mathcal{F}}$, and $\varphi^{\mathfrak{m}}$ for the coordinate at φ of a point $\mathfrak{m}$ of the product. This product, the space of all candidate assignments of values to sentences, is compact Hausdorff by Tychonov's theorem (see Engelking [Eng89, Theorem 3.2.4]), and the Henkin types form its subspace $\mathcal{H}_C^{\mathcal{F}}$, the set of assignments that satisfy the Henkin conditions. For a set T' of closed conditions on sentences of $S_C^{\mathcal{F}}$, we write $\langle\langle T' \rangle\rangle$ for the set of points that satisfy every condition of T'; in particular, $\langle\langle \emptyset \rangle\rangle$ is the whole product.

We use three results of Aubone, Dueñez and Iovino [ADI26] for the logic of this paper:

(i) A structure for $\mathcal{L}_{\mathcal{W}}$ is *canonical* when every element is the value of a closed term. The family of the values of the sentences of $S_C^{\mathcal{F}}$ in a canonical structure, its $S_C^{\mathcal{F}}$*-type*, belongs to $\mathcal{H}_C^{\mathcal{F}}$. Conversely, every member $\mathfrak{m}$ of $\mathcal{H}_C^{\mathcal{F}}$ is the $S_C^{\mathcal{F}}$-type of a canonical structure, the *term structure* of $\mathfrak{m}$ [ADI26, Lemma 3]. The elements of the term structure are the closed terms of $\mathcal{L}_{\mathcal{W}}$, on which the function symbols act by forming terms, and the value of each atomic sentence in it is the value that $\mathfrak{m}$ assigns to the sentence.

(ii) For a countable vocabulary and a countable fragment, every structure has a substructure whose sorts are countable and in which every sentence of the fragment keeps its value [ADI26, Lemma 4 and its proof]. A structure for $\mathcal{L}$ whose sorts are countable has an expansion to $\mathcal{L}_\mathcal{W}$ that names every element by a witness, and this expansion is canonical [ADI26, proof of Proposition 5].

(iii) The space $\mathcal{H}_C^\mathcal{F}$ is a G_δ subset of $\prod S_C^\mathcal{F}$ [ADI26, Proposition 6]. Every Henkin condition of Section 2.2 defines a closed set, except the upper bounds on suprema that (H5′) and (H6′) impose; each of these upper bounds defines a countable intersection of open sets. For a supremum $\sup_x \varphi(x)$, the upper bound says that its value does not exceed the supremum of the values of the sentences $\varphi(t)$, for the closed terms t of $\mathcal{L}_\mathcal{W}$ of the sort of x. For a countable join, the upper bound says that the value of the join does not exceed the supremum of the values of its members.

From (iii), Aubone, Dueñez and Iovino deduce that the space of Henkin types is Čech-complete [ADI26, Corollary 7]. The category arguments below require the same of the Henkin types that satisfy the theory. Each axiom $\psi \le r$ or $\psi \ge r$ of T defines the intersection of the open sets $\{\mathfrak{m} : \psi^\mathfrak{m} < r + \frac{1}{n}\}$ or $\{\mathfrak{m} : \psi^\mathfrak{m} > r - \frac{1}{n}\}$, for $n \ge 1$, and since T is countable, the set $\langle\!\langle T \rangle\!\rangle \cap \mathcal{H}_C^\mathcal{F}$, like $\mathcal{H}_C^\mathcal{F}$ itself, is a G_δ subset of the compact product, and hence Čech-complete [Eng89, Theorems 3.1.2 and 3.9.1]. In a Čech-complete space, by the Baire category theorem, the union of countably many nowhere dense subsets has dense complement [Eng89, Theorem 3.9.3].

Positions and plays. Every position reached in a play has its sentences in $S_C^\mathcal{F}$. With such a position $\mathfrak{p}$ we associate the set $\langle\!\langle \mathfrak{p} \rangle\!\rangle$ of the points $\mathfrak{m}$ of the product that obey its commitments and its records: $\varphi^\mathfrak{m} \in I$ for every commitment $\varphi \in I$ of $\mathfrak{p}$, and $\theta(F(\bar{w}))^\mathfrak{m} = \theta(w')^\mathfrak{m}$ for every record $F(\bar{w}) \mapsto w'$ of $\mathfrak{p}$ and every atomic formula $\theta(x)$ of $\mathcal{L}_\mathcal{W}$ whose only free variable is a variable x of the sort of w'. These equations make $F(\bar{w})$ and w' agree in every atomic formula with parameters. Since the coordinate projections are continuous, the sets $\langle\!\langle \mathfrak{p} \rangle\!\rangle$ and $\langle\!\langle T' \rangle\!\rangle$ are closed: the closed intervals of a position give a closed set, as finitely many strict conditions would give a basic open set. Satisfaction of the position in a structure requires more than agreement in atomic formulas: under an interpretation of the witnesses, its commitments must hold and every record $F(\bar{w}) \mapsto w'$ must make the term $F(\bar{w})$ and the witness w' denote the same element (see Definition 3.1). Over the Henkin types, as the following proposition shows, the difference disappears.

Proposition B.1. *Let $\mathfrak{p}$ be a position whose sentences belong to $S_C^\mathcal{F}$, and let T' be a subset of T. The set $\langle\!\langle \mathfrak{p} \rangle\!\rangle \cap \langle\!\langle T' \rangle\!\rangle \cap \mathcal{H}_C^\mathcal{F}$ is nonempty if and only if $\mathfrak{p}$ is satisfied in some model of T' under some interpretation of the witnesses.*

Proof. Let $\mathfrak{m}$ be a point of the set, and let $\mathfrak{N}$ be, by (i), a canonical structure whose $S_C^\mathcal{F}$-type is $\mathfrak{m}$. The commitments of $\mathfrak{p}$ hold in $\mathfrak{N}$, and for each record $F(\bar{w}) \mapsto w'$ of $\mathfrak{p}$, the equations defining $\langle\!\langle \mathfrak{p} \rangle\!\rangle$ make $F(\bar{w})$ and w' agree in every atomic formula with parameters. A quotient turns this agreement into the literal equality that the record requires. Call two elements of the same sort of $\mathfrak{N}$ equivalent when they agree in every atomic formula with parameters. This relation is a congruence for every function symbol, since substituting a term for a variable of an atomic formula yields an atomic formula; the quotient map therefore preserves the value of every atomic formula and, by induction on formulas, of every formula. Since $\mathfrak{m}$ lies in $\langle\!\langle T' \rangle\!\rangle$, the reduct of the quotient to $\mathcal{L}$ is a model of T'. When each witness is interpreted as the class of its value, this model satisfies $\mathfrak{p}$: its commitments hold because the quotient map preserves values, and the equations defining $\langle\!\langle \mathfrak{p} \rangle\!\rangle$ place the values of $F(\bar{w})$ and w' in one class.

Conversely, let a model $\mathfrak{M}$ of T' satisfy $\mathfrak{p}$ under an interpretation h. By (ii), applied to the expansion of $\mathfrak{M}$ by constants for the elements $h(w)$ and to the fragment generated by $\mathcal{F}$ and the sentences of $\mathfrak{p}$, the structure $\mathfrak{M}$ has a substructure $\mathfrak{M}_0$ whose sorts are countable, which contains these elements and satisfies T' and $\mathfrak{p}$ under h, the records included. The unused witnesses of each sort, infinitely many since $\mathfrak{p}$ is finite, can name every element of $\mathfrak{M}_0$. The resulting expansion is canonical, and by (i) its $S_C^{\mathcal{F}}$-type belongs to $\langle\!\langle \mathfrak{p} \rangle\!\rangle \cap \langle\!\langle T' \rangle\!\rangle \cap \mathcal{H}_C^{\mathcal{F}}$. □

The proposition is the counterpart, for positions, of the theorem that Aubone, Dueñez and Iovino call Henkin Realizability [ADI26, Proposition 5]. Taking for $\mathfrak{p}$ the empty position and $T' = T$, we see that T has a model exactly when $\langle\!\langle T \rangle\!\rangle \cap \mathcal{H}_C^{\mathcal{F}}$ is nonempty.

The proposition gives a topological meaning to both outcomes of the game. In the proof that (1) implies (2) in Theorem 3.3, Josh answers truthfully in a model of T, and every position that he reaches is satisfied in this model; by Proposition B.1, his answers keep the set $\langle\!\langle \mathfrak{p} \rangle\!\rangle \cap \langle\!\langle T \rangle\!\rangle \cap \mathcal{H}_C^{\mathcal{F}}$ nonempty at every position $\mathfrak{p}$ of the play. At a leaf of the tree of a winning strategy of Olivia, the last position $\mathfrak{p}$ is satisfied in no structure (see Corollary 3.5(2)), and Proposition B.1, with T' empty, shows that the closed set $\langle\!\langle \mathfrak{p} \rangle\!\rangle$ does not meet $\mathcal{H}_C^{\mathcal{F}}$.

Survival has the complementary meaning: a fair play that Josh survives determines a nonempty subset of $\langle\!\langle T \rangle\!\rangle \cap \mathcal{H}_C^{\mathcal{F}}$ that is closed in $\mathcal{H}_C^{\mathcal{F}}$. Let $\mathfrak{p}_0 \subseteq \mathfrak{p}_1 \subseteq \cdots$ be its positions, indexed by $i \in \mathbb{N}$, and let $\mathfrak{p}_\infty$ be their union. The subset in question is $\bigcap_i \langle\!\langle \mathfrak{p}_i \rangle\!\rangle \cap \mathcal{H}_C^{\mathcal{F}}$, and it lies in $\langle\!\langle T \rangle\!\rangle$: by fairness, every axiom of T is named in a challenge (Ax), and Josh's answer commits the sentence of the axiom to an interval on which the axiom holds. The structure $\mathfrak{M}$ built from this play as in the proof that (2) implies (1) in Theorem 3.3 is canonical, since its elements are the witnesses, each named by itself. It satisfies every commitment and every record of $\mathfrak{p}_\infty$, and by (i) its $S_C^{\mathcal{F}}$-type lies in the set, which is therefore nonempty. The play need not determine a single Henkin type. Suppose, for instance, that a predicate symbol R other than $\mathbf{1}$, with $r_R > 0$, occurs in no axiom of T; then no sentence that occurs in the play contains R. Each value of $R^{\mathfrak{M}}$ remains free in $[-r_R, r_R]$, and different choices give $S_C^{\mathcal{F}}$-types in the set that differ at $R(\bar{w})$.

Positions play the part of the conditions in Keisler's model-theoretic forcing. Keisler adjoins countably many new constants and takes as forcing conditions the finite sets of sentences of the enlarged vocabulary that are satisfiable in a given class of models, ordered by inclusion [Kei73, pp. 98 and 100, Example 1.2]. The witnesses serve as the new constants, and a position grows by commitments and records as a condition grows by sentences. By Proposition B.1, a position $\mathfrak{p}$ is satisfied in some model of T under some interpretation of the witnesses exactly when $\langle\!\langle \mathfrak{p} \rangle\!\rangle \cap \langle\!\langle T \rangle\!\rangle \cap \mathcal{H}_C^{\mathcal{F}}$ is nonempty, and these satisfiable positions correspond to the conditions for the class of models of T. A record $F(\bar{w}) \mapsto w'$ is the counterpart, in a logic without equality, of an atomic sentence $c = \tau$, where c is a new constant and τ is a closed term. Keisler requires that, for every closed term τ, each condition extend to one that contains such a sentence for some added constant c [Kei73, p. 99, clause (iv)], and the rule (Tot) meets the corresponding requirement in the game. A fair play that Josh survives yields an increasing sequence of satisfiable positions and a canonical model, as Keisler's construction yields a generic model [Kei73, p. 101, Lemma 1.4]; fairness alone, however, does not impose his requirement that every sentence of the fragment be decided.

The rules also separate universal from existential choices, as Keisler's $\forall\exists$ formulas do [Kei73, p. 107]. These formulas have the form $\forall \bar{x} \bigvee_n \exists \bar{y} (\varphi_{n1} \wedge \cdots \wedge \varphi_{nj_n})$, where n ranges over $\mathbb{N}$, each j_n is a positive integer, and the formulas φ_{ni}, for $1 \leq i \leq j_n$, belong to a fixed set of formulas of a countable fragment. Consider an axiom $\inf_{\bar{x}} \bigvee_n \sup_{\bar{y}} \theta_n(\bar{x}, \bar{y}) \geq r$ of the same shape, in which each formula θ_n is a minimum of finitely many formulas and the a priori bounds of the

formulas θ_n have a finite supremum, as the formation of the countable join requires. After Josh commits the value of the sentence of this axiom under (Ax), the rule (inf $\geq$) lets Olivia choose the witnesses for $\bar{x}$, whereas ($\bigvee \geq$) and (sup $\geq$) require Josh to choose the index n and the witnesses for $\bar{y}$, up to the tolerances that Olivia names: the universal choices are Olivia's, and the existential choices are Josh's. Keisler's general omitting types theorem asks that every condition remain satisfiable together with every instance of the given $\forall\exists$ formulas at tuples of the added constants [Kei73, p. 108]. Its counterpart in the game is the requirement that Josh answer each challenge at a position satisfied in a model of T by a position with the same property, and the proof that (1) implies (2) in Theorem 3.3 shows that he can do so whenever T has a model. When T has no model, a winning strategy of Olivia ends each play at a finite obstruction instead. The last position at each leaf holds in no structure under any interpretation of its witnesses, so that no extension of it is satisfied in a model of T or corresponds to a forcing condition for the class of models of T. The correspondence is not literal in three respects: commitments confine the values of sentences to closed intervals instead of asserting sentences, the axioms of T need not have the $\forall\exists$ shape, and Keisler's logic has equality [Kei73, p. 97], whose role in the game is taken by the records. For the strategy of Section 8.1, the contextual presentations of Section 9 picture the terminal commitments and records of the plays that reach the parity argument as vertices, identifications and intervals (see Table 2).

Closed conditions, strict conditions and compactness. The contrast between closed and strict conditions drawn in Remark 3.6 has a topological counterpart. The Henkin types that satisfy a closed condition form a closed subset of $\mathcal{H}_C^{\mathcal{F}}$, and those that satisfy a strict condition an open subset. In the compact space $\prod S_C^{\mathcal{F}}$, a family of closed sets with empty intersection has a finite subfamily with empty intersection; a family of open sets need not. For the theory $\{R > 0\} \cup \{R < \frac{1}{n} : n \geq 1\}$ of Remark 3.6, each finite part has a model, so that the open set that it defines meets the space of Henkin types, by (i) and (ii); the whole theory defines the empty set, and the tree of the winning strategy of Olivia described in Remark 3.6 for this theory branches infinitely at its root.

Closed conditions alone do not restore compactness, because the space $\mathcal{H}_C^{\mathcal{F}}$ need not be closed in the product. Were it closed and had T no model, compactness would give a finite subset T' of T with $\langle\langle T' \rangle\rangle \cap \mathcal{H}_C^{\mathcal{F}} = \emptyset$, since the axioms of T define closed sets, and by Proposition B.1 the set T' would be a finite set of axioms without models. The theory T^{01} has no model, although every finite set of its axioms has one (see Proposition 8.3); for its fragment, therefore, the space $\mathcal{H}_C^{\mathcal{F}}$ is not closed. What the category arguments use is not closedness but Čech-completeness, which $\mathcal{H}_C^{\mathcal{F}}$ owes to being a G_δ subset of the compact product. By (iii), the parts of the Henkin conditions that need not define closed sets are the upper bounds on suprema, and in the game Josh certifies such a bound by naming an instance. Under (sup $\geq$) and (inf $\leq$) the instance is a witness, chosen among the finitely many that the rules allow him: those that occur in the position and the first fresh witness. Under ($\bigvee \geq$) and ($\bigwedge \leq$) it is an index, chosen among infinitely many. Parts (3) and (4) of Corollary 3.5 record these two cases; they follow from the rules of the game alone, without the topology of $\mathcal{H}_C^{\mathcal{F}}$.

Baire category. The category argument behind Olivia's bookkeeping takes place in another space, whose points are plays. Suppose that T has a model, so that, by Theorem 3.3, Josh has a winning strategy σ'. Let $X_{\sigma'}$ be the set of infinite plays in which Josh follows σ', with the topology whose basic open sets consist of the plays that extend a given finite play. Since

the moves range over countable sets, the space $X_{\sigma'}$ is homeomorphic to a closed subspace of the Baire space $\mathbb{N}^{\mathbb{N}}$; it is therefore Polish and satisfies the Baire category theorem. For each challenge and each round, the plays in which the challenge, if it first becomes available at that round, is made then or later form a dense open subset of $X_{\sigma'}$. Since only countably many challenges occur in plays (see Section 3), the fair plays form a countable intersection of dense open subsets of $X_{\sigma'}$, which is dense by the Baire category theorem and in particular nonempty.

The bookkeeping can also be read as a strategy in a *Banach–Mazur game*. In such a game, two players alternately choose nonempty open subsets of a space, each contained in the one before, and the second player wins if the intersection of the chosen sets is contained in a given set (see Figure 14). Consider the Banach–Mazur game on $X_{\sigma'}$ in which the players choose basic open sets, that is, finite plays in which Josh follows σ', each extending the one before, and the second player wins when the infinite play that these finite plays determine is fair. The two players of this game are not Olivia and Josh: a move of either player may extend the play by many rounds, each consisting of a challenge of Olivia and an answer of Josh. Olivia's bookkeeping of Section 3, indexed by the moves of the second player rather than by the rounds of the play, supplies a winning strategy for the second player (compare Oxtoby [Oxt57]). Recall that Olivia fixes a bijection $i \mapsto (j(i), l(i))$ from $\mathbb{N}$ onto $\mathbb{N} \times \mathbb{N}$ with $j(i) \leq i$ for every $i \in \mathbb{N}$, and that at each round she enumerates the challenges then available. At the move of index k of the second player, for $k \in \mathbb{N}$, the strategy extends the current finite play by one round, in which Olivia makes the $l(k)$-th challenge of her enumeration for round $j(k)$ and Josh answers by σ'. Since each move of the second player adds a round, at least k rounds have been played, so round $j(k)$ has been reached and the challenge is still available. The finite plays therefore grow without bound and determine a unique infinite play, and the play is fair, since every pair (j, l) is of the form $(j(k), l(k))$. In the proof of Theorem 3.3, the category argument reduces to the bookkeeping, because only countably many challenges are available at each round.

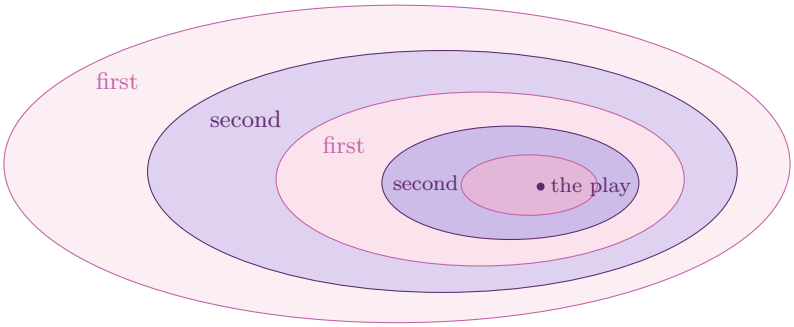


FIGURE 14. The Banach–Mazur game on the space $X_{\sigma'}$ of infinite plays in which Josh follows his winning strategy σ': nested basic open sets, chosen alternately by the first and the second player; when the second player follows the bookkeeping, the play that the sets determine is fair.

Aubone, Dueñez and Iovino prove the omitting types theorem for their logic by Morley's Baire category method, applied in the Čech-complete space of Henkin types, whose points are the types of structures built from closed terms [ADI26, Theorem 8]. To state it in the terms of this appendix, we call a *type* in a finite tuple $\bar{x}$ of variables a set $\Xi(\bar{x})$ of expressions $\varphi(\bar{x}) \geq r$, with r rational and with φ a formula of $\mathcal{F}$ whose free variables are among $\bar{x}$ [ADI26,

Section 11]. Substituting for $\bar{x}$ a tuple $\bar{t}$ of closed terms of $\mathcal{L}_{\mathcal{W}}$ of the appropriate sorts turns these expressions into closed conditions on sentences of $S_C^{\mathcal{F}}$, which form the *instance* $\Xi(\bar{t})$ of the type. For a sequence $(\Xi_n(\bar{x}_n))_{n\in\mathbb{N}}$ of types, the theorem reads as follows. Suppose that T has a model and that, for every $n \in \mathbb{N}$ and every tuple $\bar{t}$ of closed terms of $\mathcal{L}_{\mathcal{W}}$ of the appropriate sorts, the set of points of $\langle\!\langle T \rangle\!\rangle \cap \mathcal{H}_C^{\mathcal{F}}$ that satisfy the instance $\Xi_n(\bar{t})$ has empty interior in this space. Then T has a model whose sorts are countable and in which no tuple satisfies any $\Xi_n(\bar{x}_n)$. Indeed, the complements of these sets are countably many dense open subsets of the nonempty space $\langle\!\langle T \rangle\!\rangle \cap \mathcal{H}_C^{\mathcal{F}}$, and by (i) a point of their intersection, which the Baire category theorem provides, is the $S_C^{\mathcal{F}}$-type of a canonical structure whose reduct to $\mathcal{L}$ is such a model.

The hypothesis of empty interior explicates, as a statement about interiors in the Čech-complete space of Henkin types that satisfy T, the hypothesis of Keisler's general omitting types theorem recalled above, in the case in which his $\forall\exists$ formulas are taken over the whole fragment (see [ADI26, Introduction]). The interior is tested on basic open sets given by finitely many strict conditions, not on the closed sets $\langle\!\langle \mathfrak{p} \rangle\!\rangle$ attached to positions.

The two category arguments thus select different objects in different spaces: the first a fair play, the second a Henkin type that omits the prescribed types. A fair play that Josh survives determines a closed set of Henkin types, which need not reduce to a single point, and we know of no continuous map from the space of plays to the space of Henkin types that carries the dense open requirements of the one argument to those of the other. Their relation remains an analogy.

⧖ **Historical note.** The omitting types theorem goes back to Henkin, in 1954 [Hen54], and to Orey, in 1956 [Ore56]. Robinson introduced forcing into model theory in an article published in 1971 [Rob71], adapting the method that Cohen had created for set theory in 1963 [Coh63]. Barwise extended Robinson's forcing to infinitary logic and used it to give a new proof of the omitting types theorem (see Keisler [Kei73, p. 96]). Keisler described his forcing of 1973, recalled in the historical note of Section 2.2, as a reformulation, in terms of ordinary models, of the Boolean-valued models of Rasiowa and Sikorski [RS63], and he regarded its basic result, the Generic Model Theorem, as essentially due to them [Kei73, p. 97].

References

[AA66] Ichiro Amemiya and Huzihiro Araki, *A remark on Piron's paper*, Publ. Res. Inst. Math. Sci. Ser. A **2** (1966), no. 3, 423–427.

[AB11] S. Abramsky and A. Brandenburger, *The sheaf-theoretic structure of non-locality and contextuality*, New J. Phys. **13** (2011), 113036.

[Ada69] D. H. Adams, *The completion by cuts of an orthocomplemented modular lattice*, Bull. Austral. Math. Soc. **1** (1969), no. 2, 279–280.

[ADI26] Olivia Aubone, Eduardo Dueñez, and José Iovino, *Model existence and Scott's theorem for infinitary real-valued logic: a topological perspective*, arXiv:2611.02247, 2026.

[AMS21] Juan P. Aguilera, Sandra Müller, and Philipp Schlicht, *Long games and σ-projective sets*, Annals of Pure and Applied Logic **172** (2021), no. 4, 102939, arXiv:2011.04947.

[Bae52] Reinhold Baer, *Linear algebra and projective geometry*, Academic Press, New York, 1952.

[Ban56] Bernhard Banaschewski, *Hüllensysteme und Erweiterung von Quasi-Ordnungen*, Z. Math. Logik Grundlagen Math. **2** (1956), 117–130.

[BCG+22] C. Budroni, A. Cabello, O. Gühne, M. Kleinmann, and J.-Å. Larsson, *Kochen–specker contextuality*, Rev. Mod. Phys. **94** (2022), 045007.

[BDH09] Heinz H. Bauschke, Frank Deutsch, and Hein Hundal, *Characterizing arbitrarily slow convergence in the method of alternating projections*, International Transactions in Operational Research **16** (2009), no. 4, 413–425.

[Bel66] J. S. Bell, *On the problem of hidden variables in quantum mechanics*, Rev. Mod. Phys. **38** (1966), 447–452.

[Bel96] J. L. Bell, *Logical reflections on the Kochen–Specker theorem*, Perspectives on Quantum Reality (R. K. Clifton, ed.), The University of Western Ontario Series in Philosophy of Science, vol. 57, Kluwer, Dordrecht, 1996, pp. 227–235.

[Boh52] David Bohm, *A suggested interpretation of the quantum theory in terms of "hidden" variables. I and II*, Phys. Rev. **85** (1952), 166–179, 180–193.

[BvN36] Garrett Birkhoff and John von Neumann, *The logic of quantum mechanics*, Ann. of Math. (2) **37** (1936), no. 4, 823–843. MR 1503312

[BW92] L. J. Bunce and J. D. M. Wright, *The mackey–gleason problem*, Bull. Amer. Math. Soc. **26** (1992), 288–293.

[BYBHU08] Itaï Ben Yaacov, Alexander Berenstein, C. Ward Henson, and Alexander Usvyatsov, *Model theory for metric structures*, Model theory with applications to algebra and analysis. Vol. 2, London Math. Soc. Lecture Note Ser., vol. 350, Cambridge Univ. Press, Cambridge, 2008, pp. 315–427. MR MR2436146

[BYI09] Itaï Ben Yaacov and José N. Iovino, *Model theoretic forcing in analysis*, Ann. Pure Appl. Logic **158** (2009), no. 3, 163–174. MR 2500090

[BYU10] Itaï Ben Yaacov and Alexander Usvyatsov, *Continuous first order logic and local stability*, Trans. Amer. Math. Soc. **362** (2010), no. 10, 5213–5259. MR 2657678

[CEGA96] A. Cabello, J. M. Estebaranz, and G. García-Alcaine, *Bell–Kochen–Specker theorem: A proof with 18 vectors*, Phys. Lett. A **212** (1996), 183–187.

[Cha58] C. C. Chang, *Algebraic analysis of many valued logics*, Trans. Amer. Math. Soc. **88** (1958), 467–490. MR 0094302 (20 #821)

[Cha61] Chen-chung Chang, *Theory of models of infinite valued logic, I, II, III*, Notices Amer. Math. Soc. **8** (1961), 68–69.

[Chr82] E. Christensen, *Measures on projections and physical states*, Comm. Math. Phys. **86** (1982), 529–538.

[CK66] Chen-chung Chang and H. Jerome Keisler, *Continuous model theory*, Annals of Mathematics Studies, No. 58, Princeton Univ. Press, Princeton, N.J., 1966. MR MR0231708 (38 #36)

[CK00] R. Clifton and A. Kent, *Simulating quantum mechanics by non-contextual hidden variables*, Proc. R. Soc. Lond. A **456** (2000), 2101–2114.

[CKM85] R. Cooke, M. Keane, and W. Moran, *An elementary proof of gleason's theorem*, Math. Proc. Cambridge Philos. Soc. **98** (1985), 117–128.

[Coh63] Paul Cohen, *The independence of the continuum hypothesis*, Proc. Nat. Acad. Sci. U.S.A. **50** (1963), 1143–1148. MR 157890

[Ded72] Richard Dedekind, *Stetigkeit und irrationale Zahlen*, Vieweg, Braunschweig, 1872.

[DHA+13] Vincenzo D'Ambrosio, Isabelle Herbauts, Elias Amselem, Eleonora Nagali, Mohamed Bourennane, Fabio Sciarrino, and Adán Cabello, *Experimental implementation of a Kochen–Specker set of quantum tests*, Phys. Rev. X **3** (2013), 011012, arXiv:1209.1836.

[DKL15] Ehtibar N. Dzhafarov, Janne V. Kujala, and Jan-Åke Larsson, *Contextuality in three types of quantum-mechanical systems*, Found. Phys. **45** (2015), 762–782, arXiv:1411.2244.

[Dör05] A. Döring, *Kochen–specker theorem for von neumann algebras*, Int. J. Theor. Phys. **44** (2005), 139–160.

[Dvu95] Anatolij Dvurečenskij, *Frame functions and completeness of inner product spaces*, Ann. Inst. H. Poincaré Phys. Théor. **62** (1995), no. 4, 429–438.

[Eag14] Christopher J. Eagle, *Omitting types for infinitary* $[0,1]$*-valued logic*, Ann. Pure Appl. Logic **165** (2014), no. 3, 913–932. MR 3142393

[Eng89] Ryszard Engelking, *General topology*, second ed., Sigma Series in Pure Mathematics, vol. 6, Heldermann Verlag, Berlin, 1989, Translated from the Polish by the author. MR MR1039321 (91c:54001)

[Gle57] A. M. Gleason, *Measures on the closed subspaces of a Hilbert space*, J. Math. Mech. **6** (1957), 885–893.

[Göd38] Kurt Gödel, *The consistency of the axiom of choice and of the generalized continuum-hypothesis*, Proceedings of the National Academy of Sciences of the U.S.A. **24** (1938), 556–557.

[GS53] David Gale and F. M. Stewart, *Infinite games with perfect information*, Contributions to the Theory of Games, Vol. II, Annals of Mathematics Studies, no. 28, Princeton University Press, 1953, pp. 245–266.

[GZ88] C. D. Godsil and J. Zaks, *Colouring the sphere*, 1988, Univ. Waterloo research report CORR 88-12 (1988); arXiv:1201.0486.

[Háj98] Petr Hájek, *Metamathematics of fuzzy logic*, Trends in Logic—Studia Logica Library, vol. 4, Kluwer Academic Publishers, Dordrecht, 1998. MR 1900263 (2003c:03048)

[Ham03] J. Hamhalter, *Quantum measure theory*, Fundamental Theories of Physics, vol. 134, Kluwer Academic Publishers, Dordrecht, 2003.

[Har91] John Harding, *Orthomodular lattices whose MacNeille completions are not orthomodular*, Order **8** (1991), no. 1, 93–103.

[Hen49] Leon Henkin, *The completeness of the first-order functional calculus*, J. Symbolic Logic **14** (1949), 159–166. MR 33781

[Hen54] ———, *A generalization of the concept of ω-consistency*, J. Symbolic Logic **19** (1954), 183–196. MR 63324

[Hen76] C. W. Henson, *Nonstandard hulls of Banach spaces*, Israel J. Math. **25** (1976), no. 1-2, 108–144.

[HI02] C. Ward Henson and José N. Iovino, *Ultraproducts in analysis*, Analysis and logic (Mons, 1997), London Math. Soc. Lecture Note Ser., vol. 262, Cambridge Univ. Press, Cambridge, 2002, pp. 1–110. MR 1 967 834

[Hil99] David Hilbert, *Grundlagen der Geometrie*, Festschrift zur Feier der Enthüllung des Gauss–Weber-Denkmals in Göttingen, Teubner, Leipzig, 1899.

[Hol95] Samuel S. Holland, Jr., *Orthomodularity in infinite dimensions; a theorem of M. Solèr*, Bull. Amer. Math. Soc. (N.S.) **32** (1995), no. 2, 205–234. MR 1307904

[Hoo78] D. N. Hoover, *Probability logic*, Ann. Math. Logic **14** (1978), no. 3, 287–313.

[HP87] Jan Hamhalter and Pavel Pták, *A completeness criterion for inner product spaces*, Bull. London Math. Soc. **19** (1987), no. 3, 259–263.

[HP04] E. Hrushovski and I. Pitowsky, *Generalizations of kochen and specker's theorem and the effectiveness of gleason's theorem*, Stud. Hist. Philos. Mod. Phys. **35** (2004), 177–194.

[Hus37] Kôdi Husimi, *Studies on the foundation of quantum mechanics. I*, Proc. Phys.-Math. Soc. Japan (3) **19** (1937), 766–789.

[Iov94] José N. Iovino, *Stable theories in functional analysis*, Ph.D. thesis, University of Illinois at Urbana-Champaign, 1994.

[Iov99a] ———, *Stable Banach spaces and Banach space structures. I. Fundamentals*, Models, algebras, and proofs (Bogotá, 1995), Dekker, New York, 1999, pp. 77–95. MR 2000h:03064a

[Iov99b] ———, *Stable Banach spaces and Banach space structures. II. Forking and compact topologies*, Models, algebras, and proofs (Bogotá, 1995), Dekker, New York, 1999, pp. 97–117. MR 2000h:03064b

[Kar64] Carol R. Karp, *Languages with expressions of infinite length*, North-Holland, Amsterdam, 1964.

[KDL15] Janne V. Kujala, Ehtibar N. Dzhafarov, and Jan-Åke Larsson, *Necessary and sufficient conditions for an extended noncontextuality in a broad class of quantum mechanical systems*, Phys. Rev. Lett. **115** (2015), 150401.

[Kec95] Alexander S. Kechris, *Classical descriptive set theory*, Graduate Texts in Mathematics, vol. 156, Springer-Verlag, New York, 1995. MR 1321597

[Kei71] H. Jerome Keisler, *Model theory for infinitary logic. Logic with countable conjunctions and finite quantifiers*, Studies in Logic and the Foundations of Mathematics, vol. Vol. 62, North-Holland Publishing Co., Amsterdam-London, 1971. MR 344115

[Kei73] ———, *Forcing and the omitting types theorem*, Studies in model theory, Math. Assoc. Amer., Buffalo, N.Y., 1973, pp. 96–133. MAA Studies in Math., Vol. 8. MR 49 #2340

[Kei85] ———, *Probability quantifiers*, Model-Theoretic Logics (J. Barwise and S. Feferman, eds.), Perspectives in Mathematical Logic, Springer-Verlag, New York, 1985, pp. 509–556.

[Kei23] H. Jerome Keisler, *Model theory for real-valued structures*, Beyond First-Order Model Theory, Vol. 2 (José Iovino, ed.), CRC Press, 2023, arXiv:2005.11851, pp. 43–101.

[Kel80] Hans Arwed Keller, *Ein nicht-klassischer Hilbertscher Raum*, Math. Z. **172** (1980), no. 1, 41–49.

[Ken99] A. Kent, *Noncontextual hidden variables and physical measurements*, Phys. Rev. Lett. **83** (1999), 3755–3757.

[Ker94] M. Kernaghan, *Bell–Kochen–Specker theorem for 20 vectors*, J. Phys. A: Math. Gen. **27** (1994), L829–L830.

[KM01] H. Jerome Keisler and Arnold W. Miller, *Categoricity without equality*, Fund. Math. **170** (2001), no. 1-2, 87–106, Dedicated to the memory of Jerzy Łoś. MR 1881370 (2002k:03054)

[Kőn27] D. Kőnig, *Über eine Schlußweise aus dem Endlichen ins Unendliche*, Acta Sci. Math. (Szeged) **3** (1927), 121–130.

[KPS23] M. Kirchweger, T. Peitl, and S. Szeider, *Co-certificate learning with sat modulo symmetries*, Proc. IJCAI 2023, 2023, arXiv:2306.10427.

[KS67] Simon Kochen and E. P. Specker, *The problem of hidden variables in quantum mechanics*, J. Math. Mech. **17** (1967), 59–87. MR 219280

[KS15] Ravi Kunjwal and Robert W. Spekkens, *From the Kochen–Specker theorem to noncontextuality inequalities without assuming determinism*, Phys. Rev. Lett. **115** (2015), 110403.

[KW88] Selahattin Kayalar and Howard L. Weinert, *Error bounds for the method of alternating projections*, Mathematics of Control, Signals, and Systems **1** (1988), 43–59.

[Lar02] Jan-Åke Larsson, *A Kochen–Specker inequality*, Europhys. Lett. **58** (2002), no. 6, 799–805, arXiv:quant-ph/0006134.

[Lar04] Paul B. Larson, *The stationary tower: Notes on a course by W. Hugh Woodin*, University Lecture Series, vol. 32, American Mathematical Society, Providence, RI, 2004.

[LBG24] Z. Li, C. Bright, and V. Ganesh, *A SAT solver and computer algebra attack on the minimum Kochen–Specker problem*, Proc. IJCAI 2024, 2024, arXiv:2306.13319.

[LE65] E. G. K. Lopez-Escobar, *An interpolation theorem for denumerably long formulas*, Fundamenta Mathematicae **57** (1965), 253–272.

[Lus25] Nicolas Lusin, *Sur les ensembles projectifs de M. Henri Lebesgue*, Comptes Rendus de l'Académie des Sciences, Paris **180** (1925), 1572–1574.

[Lö15] Leopold Löwenheim, *Über möglichkeiten im relativkalkül*, Mathematische Annalen **76** (1915), 447–470.

[Mac37] H. M. MacNeille, *Partially ordered sets*, Trans. Amer. Math. Soc. **42** (1937), no. 3, 416–460.

[Mac57] George W. Mackey, *Quantum mechanics and Hilbert space*, Amer. Math. Monthly **64** (1957), no. 8, part II, 45–57.

[Mak69a] M. Makkai, *An application of a method of Smullyan to logics on admissible sets*, Bull. Acad. Polon. Sci. Sér. Sci. Math. Astronom. Phys. **17** (1969), 341–346.

[Mak69b] ———, *On the model theory of denumerably long formulas with finite strings of quantifiers*, J. Symbolic Logic **34** (1969), no. 3, 437–459.

[Mau81] R. Daniel Mauldin (ed.), *The Scottish Book: Mathematics from the Scottish Café*, Birkhäuser, Boston, 1981.

[Mer90] N. David Mermin, *Simple unified form for the major no-hidden-variables theorems*, Phys. Rev. Lett. **65** (1990), no. 27, 3373–3376.

[Mey99] D. A. Meyer, *Finite precision measurement nullifies the Kochen–Specker theorem*, Phys. Rev. Lett. **83** (1999), 3751–3754.

[MM70] F. Maeda and S. Maeda, *Theory of symmetric lattices*, Die Grundlehren der mathematischen Wissenschaften, vol. Band 173, Springer-Verlag, New York-Berlin, 1970. MR 282889

[Mor73] Ronald P. Morash, *Angle bisection and orthoautomorphisms in Hilbert lattices*, Canadian J. Math. **25** (1973), 261–272. MR 313143

[Mor74] Michael Morley, *Applications of topology to $L_{\omega 1\omega}$*, Proceedings of the Tarski Symposium (Proc. Sympos. Pure Math., Vol. XXV, Univ. California, Berkeley, Calif., 1971) (Providence, R.I.), Amer. Math. Soc., 1974, pp. 233–240. MR 0373879 (51 #10079)

[Mos09] Yiannis N. Moschovakis, *Descriptive set theory*, second ed., Mathematical Surveys and Monographs, vol. 155, American Mathematical Society, Providence, RI, 2009.

[MŚ64] Jan Mycielski and Stanisław Świerczkowski, *On the Lebesgue measurability and the axiom of determinateness*, Fundamenta Mathematicae **54** (1964), 67–71.

[MS89] Donald A. Martin and John R. Steel, *A proof of projective determinacy*, Journal of the American Mathematical Society **2** (1989), no. 1, 71–125.

[Ore56] Steven Orey, *On ω-consistency and related properties*, J. Symbolic Logic **21** (1956), 246–252. MR 82936

[Oxt57] John C. Oxtoby, *The Banach–Mazur game and Banach category theorem*, Contributions to the Theory of Games, Vol. III, Annals of Mathematics Studies, vol. 39, Princeton University Press, 1957, pp. 159–163.

[Per90] Asher Peres, *Incompatible results of quantum measurements*, Phys. Lett. A **151** (1990), no. 3-4, 107–108.

[Per91] A. Peres, *Two simple proofs of the Kochen–Specker theorem*, J. Phys. A: Math. Gen. **24** (1991), L175–L178.

[Pir64] C. Piron, *Axiomatique quantique*, Helv. Phys. Acta **37** (1964), 439–468.

[Pit98] I. Pitowsky, *Infinite and finite gleason's theorems and the logic of indeterminacy*, J. Math. Phys. **39** (1998), 218–228.

[Pre95] Alexander Prestel, *On Solèr's characterization of Hilbert spaces*, Manuscripta Math. **86** (1995), no. 2, 225–238. MR 1317747

[Réd07] Miklós Rédei, *The birth of quantum logic*, Hist. Philos. Logic **28** (2007), no. 2, 107–122.

[Rob71] Abraham Robinson, *Forcing in model theory*, Symposia Mathematica, Vol. V (INDAM, Rome, 1969/70), Academic Press, London-New York, 1971, pp. 69–82. MR 278925

[RS63] Helena Rasiowa and Roman Sikorski, *The mathematics of metamathematics*, Monografie Matematyczne, vol. 41, Państwowe Wydawnictwo Naukowe, Warsaw, 1963.

[SBZ01] Christoph Simon, Časlav Brukner, and Anton Zeilinger, *Hidden-variable theorems for real experiments*, Phys. Rev. Lett. **86** (2001), 4427, arXiv:quant-ph/0006043.

[Sch56] Jürgen Schmidt, *Zur Kennzeichnung der Dedekind-MacNeilleschen Hülle einer geordneten Menge*, Arch. Math. **7** (1956), 241–249.

[Sie22] Wacław Sierpiński, *Sur les fonctions d'ensemble additives et continues*, Fund. Math. **3** (1922), 240–246.

[Sie25] ———, *Sur une classe d'ensembles*, Fundamenta Mathematicae **7** (1925), 237–243.

[Sko70] Thoralf Skolem, *Selected works in logic*, Universitetsforlaget, Oslo, 1970, Edited by Jens Erik Fenstad. MR 285342

[Sol95] M. P. Solèr, *Characterization of Hilbert spaces by orthomodular spaces*, Comm. Algebra **23** (1995), no. 1, 219–243. MR 1311786

[Spe60] Ernst Specker, *Die Logik nicht gleichzeitig entscheidbarer Aussagen*, Dialectica **14** (1960), 239–246.

[Spe05] R. W. Spekkens, *Contextuality for preparations, transformations, and unsharp measurements*, Phys. Rev. A **71** (2005), 052108, arXiv:quant-ph/0406166.

[UW16] Sander Uijlen and Bas Westerbaan, *A Kochen–Specker system has at least 22 vectors*, New Gener. Comput. **34** (2016), 3–23.

[Vää11] Jouko Väänänen, *Models and games*, Cambridge Studies in Advanced Mathematics, vol. 132, Cambridge University Press, Cambridge, 2011.

[vN49] John von Neumann, *On rings of operators. Reduction theory*, Ann. of Math. (2) **50** (1949), 401–485.

[vN50] ———, *Functional operators. Vol. II: The geometry of orthogonal spaces*, Annals of Mathematics Studies, no. 22, Princeton University Press, Princeton, N.J., 1950, Reprint of mimeographed lecture notes first distributed in 1933.

[VY08] Oswald Veblen and John Wesley Young, *A set of assumptions for projective geometry*, Amer. J. Math. **30** (1908), 347–380.

[VY10] ———, *Projective geometry, vol. I*, Ginn and Co., Boston, 1910.

[XCG20] Zhen-Peng Xu, Jing-Ling Chen, and Otfried Gühne, *Proof of the Peres conjecture for contextuality*, Phys. Rev. Lett. **124** (2020), 230401.

[Yea83] F. J. Yeadon, *Measures on projections in w^*-algebras of type ii_1*, Bull. London Math. Soc. **15** (1983), 139–145.

Department of Mathematics, UT San Antonio, One UTSA Circle, San Antonio, TX 78249, USA

Email address: olivia.aubone@utsa.edu

Department of Mathematics, UT San Antonio, One UTSA Circle, San Antonio, TX 78249, USA

Email address: jose.iovino@utsa.edu