

A COUNTEREXAMPLE TO VAUGHT’S CONJECTURE FOR $\mathcal{L}_{\omega_1, \omega}$

NATHANAEL L. ACKERMAN, CAMERON E. FREER, AND ROBIN W. KNIGHT

ABSTRACT. We construct a sentence of $\mathcal{L}_{\omega_1, \omega}$ having exactly $\aleph_1$ many countable models up to isomorphism, but no perfect set of non-isomorphic models. This disproves Vaught’s conjecture for infinitary logic. The construction follows a framework devised by Robin Knight, and the proof has been formally verified in Lean.

CONTENTS

1. Introduction	1
1.1. Previous Results and Approaches	2
1.2. History of Our Proof	3
1.3. Outline	4
1.4. Notation	4
2. Knight Systems	4
2.1. Geometry Languages	4
2.2. Invariant Diagrams and Invariant Systems	6
2.3. Density	7
2.4. Back and Forth Systems	8
2.5. Weak Knight Systems	9
2.6. Knight Systems and Knight Classes	10
3. Toy Example	11
3.1. Labeled Geometric Trees	11
3.2. Knight System: Attempt 1	13
3.3. Knight System: Attempt 2	15
4. The Counterexample	16
4.1. Plans and Templates	17
4.2. Knight Systems	19
4.3. Legal Templates	21
References	25

1. INTRODUCTION

The beginning of modern model theory can be traced back to a paper of Vaught titled “Denumerable models of complete theories” [Vau61]. That paper introduced many of the fundamental concepts and ideas used in model theory, and it ended with the following question which gave rise to the conjecture that bears his name:

Question 1.1 ([Vau61, p. 320]). *Can it be proved, without the use of the continuum hypothesis, that there exists a complete theory having exactly $\aleph_1$ non-isomorphic denumerable models?*

2020 *Mathematics Subject Classification*. Primary: 03C15; Secondary: 03C75, 03E15.

Key words and phrases. Vaught’s conjecture, Knight systems, number of models, infinitary logic.

The statement that all countable first-order theories either have countably many or continuum many countable models (up to isomorphism) has since been known as “Vaught’s conjecture”.

The most significant advance in the study of Vaught’s conjecture of the next several decades was Morley’s paper “The number of countable models” [Mor70]. In this paper Morley made two important advances. First, he showed that if Vaught’s conjecture were to fail, then any counterexample must be a scattered theory with precisely $\aleph_1$ many models. Second, he showed that these results were true not just of first-order logic but also of sentences of the infinitary logic $\mathcal{L}_{\omega_1, \omega}$.

In light of the first advance, Vaught’s conjecture is usually taken today to mean the “absolute” version, namely the statement that every countable first-order theory with an uncountable collection of models (up to isomorphism) has a (non-empty) perfect set of non-isomorphic models. In particular, this formulation, unlike the original, is absolute between models of set theory and does not depend on whether the continuum hypothesis holds.

Morley’s results for $\mathcal{L}_{\omega_1, \omega}$ not only suggested a natural generalization of Vaught’s conjecture for sentences of $\mathcal{L}_{\omega_1, \omega}$, but also opened the door for a much wider generalization. Several years earlier, Lopez-Escobar [LE65] showed that there is a close connection between the countable model theory of $\mathcal{L}_{\omega_1, \omega}$ and the descriptive set theory of continuous actions of S_∞ on Polish spaces. This implies that Vaught’s conjecture for $\mathcal{L}_{\omega_1, \omega}$ can be seen as the S_∞ case of what is known as the topological Vaught conjecture.

Conjecture 1.2 (Topological Vaught Conjecture). *If a Polish group G acts continuously on a Polish space X , then either there are countably many orbits, or there is a (non-empty) perfect subset of X whose distinct points lie in distinct orbits.*

In this way, the topological Vaught conjecture for the Polish group S_∞ can be seen as one of the oldest problems in both model theory and descriptive set theory.

1.1. Previous Results and Approaches. Over the years, Vaught’s conjecture has been attacked from many different directions, but two are especially worth mentioning. The first approach encompasses a wide body of research showing that in many well-behaved situations Vaught’s conjecture does indeed hold. These can, roughly, be grouped into four classes.

First, there are those specific types of structures on which Vaught’s conjecture has been shown to hold. These include any theory of Boolean algebras [Ive91], modules over a serial ring [Pun00], varieties [HSV94], trees [Ste78], and theories with a definable discrete linear order [Tan24].

Second, there are abstract properties of the theories which imply Vaught’s conjecture. These include ω -stable theories [SHM84], superstable theories of finite rank [Bue08], o-minimal theories [May88], almost chainable theories [Kur21] and sentences of $\mathcal{L}_{\omega_1, \omega}$ in which the equality symbol is not used [Ack15].

Third is work showing that certain higher computability bounds on the complexity of models imply Vaught’s conjecture. One of the most important papers in this vein is [Sac83]. There Sacks showed that if either no model of a given theory has its maximum allowed Scott rank, or if no Scott rank has multiple models of the theory, then Vaught’s conjecture holds of that theory. More recently, in [Mon13] it was shown (assuming ZFC + Projective Determinacy) that a theory with uncountably many countable models up to isomorphism is a counterexample if and only if it satisfies “hyperarithmetic-is-recursive on a cone”.

The fourth class of examples where Vaught’s conjecture is known to hold deals with the study of Polish groups and the topological Vaught conjecture. In [HS99] it was shown that the topological Vaught conjecture held for all continuous actions of nilpotent Polish groups on Polish spaces. In [Sam94] it was shown, among other things, that Vaught’s conjecture holds for all abelian Polish groups.

The second approach to attacking Vaught’s conjecture consisted in studying consequences of the existence of a counterexample. For example, in [Bec13] it was shown that if Vaught’s conjecture fails, then there must be a sentence of $\mathcal{L}_{\omega_1, \omega}$ and a $b \in 2^\omega$ such that the class of all b -computable models is hyperarithmetic in b , but exactly one isomorphism class has an index set that is not hyperarithmetic in b .

In [Bec94] it was shown that if a group has an action that is a counterexample to the topological Vaught conjecture then there must be a “minimal” such action, i.e., one where for any invariant Borel set C , either C or its complement contains countably many orbits. The paper [Hjo07] contained the striking result that if there is a counterexample to Vaught’s conjecture for $\mathcal{L}_{\omega_1, \omega}$, then there must be a counterexample for $\mathcal{L}_{\omega_1, \omega}$ which has no models of size $\aleph_2$.

Hjorth’s tour-de-force [Hjo01] established that if G is a Polish group which does not satisfy the topological Vaught conjecture then there must be a closed subgroup of G which has S_∞ as a continuous homomorphic image.

Finally, for an informative survey of various different approaches which have been tried, see [BFKL16].

1.2. History of Our Proof. Our proof that there is a counterexample to Vaught’s conjecture for $\mathcal{L}_{\omega_1, \omega}$ arose in a very unusual way. It began back in 2002 when Knight announced he had proved that there is a first-order counterexample to Vaught’s conjecture [Kni02]. Given the importance of the result, the community spent a great deal of effort trying to validate Knight’s paper, including a conference at Notre Dame dedicated to the paper, which both Knight and Ackerman attended. However, despite these efforts, Knight’s argument did not achieve general acceptance as a correct proof, although no definitive flaw in the construction had been identified.

Over the next several years, Knight attempted several new formulations of his approach, with the goal of making it more understandable by the community. During that time Ackerman read, and gave feedback on, several of Knight’s drafts, including some which were not made public.

Then, in February 2026, Ackerman and Freer approached Knight with the idea to use AI coding agents to help formalize his proof in Lean. Knight agreed and gave Ackerman and Freer a draft of a new, clearer version of the infinitary argument which he had been working on but had not yet made public. While Knight had not thoroughly vetted this proof for errors, it was, at its core, a streamlined presentation of his earlier approach.

Over the following 8 months, Freer worked extensively with coding agents (Claude Code and Codex with models ranging from Opus 4.5 through 5.5, Fable 5 and 5.1, and GPT 5.2 through 6) and occasional help from GPT Pro (5.2 through 6) to formalize the infinitary argument from Knight’s recent working draft. This involved using an agentic coding “skill” `lean4-skills` [Fre25] that Freer had developed to streamline and improve the process of using Lean, as well as the Lean LSP MCP server [Dre25]. During this time, Freer periodically consulted with Ackerman as difficulties arose.

Eventually this process led to a formally verified counterexample to Vaught’s conjecture for infinitary logic, with the entire argument certified in Lean. The formal proof makes use of Freer’s `infinitary-logic` Lean library [Fre26] for standard results in the model theory of $\mathcal{L}_{\omega_1, \omega}$ and descriptive set theory. A version of the complete formal proof has been checked by Comparator and is available on the Palomar Registry [AFK26a]. Another version of the formal proof, available on GitHub [AFK26b], has a cleaner API so as to provide reusable building blocks for further refinements and extensions. It follows the present exposition somewhat more closely than the first complete Lean proof, although its proof organization still differs in some places, and not every line of this exposition has been formally verified.

The process of producing the counterexample turned out to be less straightforward than typical formalization projects. Knight’s working draft presented a very complicated combinatorial object along with a strategy for deriving a counterexample from its various structural and extension properties. While the properties that Knight described were essentially sufficient to yield a counterexample, substantial revisions to the combinatorial object were required in order for it to satisfy these conditions. Closing this gap involved a great deal of human mathematical direction and AI-assisted exploration. Near the end, SMT and Horn clause solvers were used to investigate finite configurations and test proposed lemmas, though the final Lean proof does not involve trusting these solvers, and subsequently the resulting complicated objects were substantially simplified.

Once a counterexample was found, Freer helped guide the AI to further streamline the argument while Ackerman wrote the first draft of this paper. Our goal is for the presentation in this paper to

describe abstractions that highlight the key ideas present in all of Knight's papers while also outlining the overall path to a formally verified proof. Many details are elided, especially near the end of the argument; we plan to describe the argument more fully in a forthcoming manuscript.

1.3. Outline. After this outline and some notation, the rest of this paper has three main sections. First, in Section 2, we introduce “Knight systems”, a framework for showing that a class of countable structures has exactly $\aleph_1$ many isomorphism types and no perfect set of non-isomorphic structures. While the *specific* framework of Knight systems is introduced here for the first time by Ackerman, it abstracts the central structural ideas developed throughout Knight's earlier work, going back to his 2002 paper [Kni02].

Second, in Section 3, we introduce a “toy example” of a Knight system built from trees. While the class of structures will not end up being the set of models of a sentence of $\mathcal{L}_{\omega_1, \omega}$, it is instructive to see how one might try to build a Knight system from trees, and many of the features we see in trees will occur in the eventual counterexample. As such, this toy example is a useful framework to keep in mind while reading through the much more complicated section on the counterexample.

Third, in Section 4, we introduce the counterexample in three parts. First, in Section 4.1, we introduce the basic components of the counterexample, culminating in the notion of a template system. Second, in Section 4.2, we discuss how one would build a weak Knight system out of a template system. Finally, in Section 4.3, we introduce the combinatorial notion of “legal templates”. This is defined specifically so that the weak Knight system built from them would be a Knight system and hence witness that they describe a counterexample. We then derive the main theorem, that the weak Knight system constructed in Section 4.2 is in fact a Knight system, from three substantial propositions. These propositions were the main aspects which required a large amount of guided AI work. They have been formally verified in Lean, but their proofs are omitted here.

1.4. Notation. We now give some conventions and notation which are specific to this paper. All languages will be relational and all relation symbols have arity at least 1. For a language $\mathfrak{L}$, let $\mathfrak{L}^n \subseteq \mathfrak{L}$ be the language consisting of only those relation symbols in $\mathfrak{L}$ of arity n . All structures will be countable with non-empty domains. We will use script letters for structures with the corresponding Roman letters for their domain, e.g., $\mathcal{M}$ has underlying set M . Bold letters will represent tuples of elements or tuples of variables. We will write $\mathbf{a} \subseteq \mathcal{M}$ to mean that all elements of the tuple $\mathbf{a}$ are in M . For $n \in \omega$, let $[n] = \{0, 1, \dots, n-1\}$.

If E is a symmetric relation on a set X we say that a subset $X_0 \subseteq X$ is an *E -antichain* if $(\forall x, y \in X_0) (x \neq y) \rightarrow \neg E(x, y)$.

If $\mathfrak{L}$ is a countable language let $\text{Str}_{\mathfrak{L}}(\omega)$ be the collection of $\mathfrak{L}$ -structures with underlying set ω given the standard topology. We say that a non-empty closed subset of $\text{Str}_{\mathfrak{L}}(\omega)$ is a **perfect set** of models if it has no isolated points.

2. KNIGHT SYSTEMS

In this section we introduce the key abstractions of this paper, culminating in the notion of a Knight system. We then show that any class of structures which is the base of a Knight system must have $\aleph_1$ many structures up to isomorphism as well as no perfect set of non-isomorphic structures.

2.1. Geometry Languages. An important aspect of the construction is a background finitary geometry on each structure, and we will only attempt to characterize the structure of the finite sets which are closed in this geometry.

Definition 2.1. Suppose $\mathfrak{L}_a$ is a language where, for each $n \geq 1$, the language $\mathfrak{L}_a^n$ has a unique relation symbol P_n . We define the **theory of $\mathfrak{L}_a$ -geometry** to be the countable theory $\text{Th}_{\text{Geo}}(\mathfrak{L}_a) \subseteq \mathcal{L}_{\omega_1, \omega}(\mathfrak{L}_a)$ consisting of the following sentences.

- For $n \geq 1$,

$$(\forall x_0, \dots, x_{n-1}) P_n(x_0, \dots, x_{n-1}) \rightarrow \bigwedge_{i < j \in [n]} (x_i \neq x_j)$$

and if $\sigma: [n] \rightarrow [n]$ is a bijection then

$$(\forall x_0, \dots, x_{n-1}) P_n(x_0, \dots, x_{n-1}) \leftrightarrow P_n(x_{\sigma(0)}, \dots, x_{\sigma(n-1)}).$$

- For all $n \geq 1$,

$$(\forall x_0, \dots, x_{n-1}) \bigwedge_{i < j \in [n]} (x_i \neq x_j) \rightarrow \bigvee_{k \in \omega} (\exists x_n, \dots, x_{n+k-1}) P_{n+k}(x_0, \dots, x_{n+k-1}).$$

- For $n_0, n_1, k \geq 1$ if $P_{n_0}(x_0, \dots, x_{n_0-1})$ and $P_{n_1}(y_0, \dots, y_{n_1-1})$ and $z_0, \dots, z_{k-1}$ is an enumeration of $\{x_i\}_{i \in [n_0]} \cap \{y_i\}_{i \in [n_1]}$ then $P_k(z_0, \dots, z_{k-1})$ holds.

If $\mathfrak{L}_a \subseteq \mathfrak{L}$ and $\mathcal{M}$ is an $\mathfrak{L}$ -structure which satisfies $\text{Th}_{\text{Geo}}(\mathfrak{L}_a)$ then we say $\mathcal{M}$ has a $\mathfrak{L}_a$ -**geometry** and we call a tuple satisfying a relation in $\mathfrak{L}_a$ **closed**.

An $\mathfrak{L}_a$ -structure $\mathcal{M}$ satisfies $\text{Th}_{\text{Geo}}(\mathfrak{L}_a)$ if there is a locally finite closure relation $\text{cl}_{\mathcal{M}}$ on M such that $\{\mathbf{a} : \mathcal{M} \models P_n(\mathbf{a})\}$ is precisely the collection of enumerations of non-empty closed sets of size n .

For our purposes, it is not enough to simply define the closed sets. We also need to define, for each tuple enumerating a closed set, an “invariant” of the tuple. This invariant is meant to capture all relevant information about the tuple.

Definition 2.2. Suppose that

- $\mathfrak{L}_a, \mathfrak{L}_b, \mathfrak{L}_c$ are languages with $\mathfrak{L}_c \cap \mathfrak{L}_b = \emptyset$ and $\mathfrak{L}_a \subseteq \mathfrak{L}_c$, and
- for each $n \geq 1$, we have $|\mathfrak{L}_a^n| = 1$ and $|\mathfrak{L}_b^n| > 0$.

We define the theory of $(\mathfrak{L}_a, \mathfrak{L}_b, \mathfrak{L}_c)$ -**structured geometry**, denoted by $\text{Th}_{\text{SGeo}}(\mathfrak{L}_a, \mathfrak{L}_b, \mathfrak{L}_c)$, to consist of the following collections of sentences.

- $\text{Th}_{\text{Geo}}(\mathfrak{L}_a)$.
- For all $n \geq 1$ and $P_n \in \mathfrak{L}_a^n$,

$$(\forall x_0, \dots, x_{n-1}) P_n(x_0, \dots, x_{n-1}) \leftrightarrow \bigvee_{Q_n \in \mathfrak{L}_b^n} Q_n(x_0, \dots, x_{n-1}).$$

- For all $n \geq 1$ and distinct $Q_n, Q'_n \in \mathfrak{L}_b^n$,

$$(\forall x_0, \dots, x_{n-1}) Q_n(x_0, \dots, x_{n-1}) \rightarrow \neg Q'_n(x_0, \dots, x_{n-1}).$$

- For any $k, n \geq 1$, any $Q \in \mathfrak{L}_b^n$, any $R \in \mathfrak{L}_b^k \cup \mathfrak{L}_c^k$, and any sequence $\langle i_j \rangle_{j \in [k]} \subseteq [n]$,

$$(\forall x_0, \dots, x_{n-1}) (\forall y_0, \dots, y_{k-1}) [Q(x_0, \dots, x_{n-1}) \wedge Q(y_0, \dots, y_{k-1})] \\ \rightarrow [R(x_{i_0}, \dots, x_{i_{k-1}}) \leftrightarrow R(y_{i_0}, \dots, y_{i_{k-1}})].$$

If $\mathfrak{L}_a, \mathfrak{L}_b, \mathfrak{L}_c \subseteq \mathfrak{L}$ and $\mathcal{M}$ is an $\mathfrak{L}$ -structure which satisfies $\text{Th}_{\text{SGeo}}(\mathfrak{L}_a, \mathfrak{L}_b, \mathfrak{L}_c)$ then we say that $\mathcal{M}$ has an $(\mathfrak{L}_a, \mathfrak{L}_b, \mathfrak{L}_c)$ -**structured geometry**.

We call elements of $\mathfrak{L}_b$ **invariants**. If $\mathcal{M} \models \text{Th}_{\text{SGeo}}(\mathfrak{L}_a, \mathfrak{L}_b, \mathfrak{L}_c)$ and $\mathcal{M} \models Q(\mathbf{a})$ where $Q \in \mathfrak{L}_b$ then we call Q the **invariant** of $\mathbf{a}$. We also refer to $\mathcal{M}$ as a **lift** of $\mathcal{M}|_{\mathfrak{L}_c}$.

So $\mathcal{M}$ has an $(\mathfrak{L}_a, \mathfrak{L}_b, \mathfrak{L}_c)$ -structured geometry if (a) $\mathcal{M}$ has an $\mathfrak{L}_a$ -geometry, (b, c) every enumeration of a closed set realizes a unique invariant in $\mathfrak{L}_b$, and (d) the relation symbols in $\mathfrak{L}_b$ completely determine the quantifier-free $(\mathfrak{L}_b \cup \mathfrak{L}_c)$ -structure of their arguments. A useful intuition is that the relation symbols in $\mathfrak{L}_a$ define a finitary geometry and the relation symbols in $\mathfrak{L}_b$ provide “invariants” for each finite closed set.

Convention 2.3. For the rest of Section 2 we fix languages $\mathfrak{L}_a, \mathfrak{L}_b, \mathfrak{L}_c$ with $\mathfrak{L}_a \subseteq \mathfrak{L}_c$, $\mathfrak{L}_b \cap \mathfrak{L}_c = \emptyset$ and such that for all $n \geq 1$, we have $\mathfrak{L}_a^n = \{P_n\}$ and $|\mathfrak{L}_b^n| > 0$ and $|\mathfrak{L}_c^n| \leq \omega$.

2.2. Invariant Diagrams and Invariant Systems. A key feature of structured geometries is that an invariant determines which atomic formulas hold on its arguments. In this way an invariant of a tuple determines the structure of a tuple. We will be interested in classes where the structure of a specific invariant is the same across all models in the class.

Definition 2.4. An **invariant diagram** is a collection $\mathcal{I}$ of pairs of formulas satisfying the following.

- (a) If $(P(x_0, \dots, x_{n-1}), \psi(y_0, \dots, y_{k-1})) \in \mathcal{I}$ then $\{y_i\}_{i \in [k]} \subseteq \{x_i\}_{i \in [n]}$ and $P \in \mathfrak{L}_b$, and either $\psi(y_0, \dots, y_{k-1})$ is an atomic formula or the negation of an atomic formula.
- (b) For every $Q \in \mathfrak{L}_b^n$, $R \in \mathfrak{L}_b^k \cup \mathfrak{L}_c^k$ and $\langle j_i \rangle_{i \in [k]} \subseteq [n]$ precisely one of the following holds:
 - $(Q(x_0, \dots, x_{n-1}), R(x_{j_0}, \dots, x_{j_{k-1}})) \in \mathcal{I}$, or
 - $(Q(x_0, \dots, x_{n-1}), \neg R(x_{j_0}, \dots, x_{j_{k-1}})) \in \mathcal{I}$.
- (c) For each $Q \in \mathfrak{L}_b$ there is an $(\mathfrak{L}_b \cup \mathfrak{L}_c)$ -structure $\mathcal{M}$ where
 - $\mathcal{M} \models \text{Th}_{\text{SGeo}}(\mathfrak{L}_a, \mathfrak{L}_b, \mathfrak{L}_c)$,
 - $\mathcal{M} \models (\exists \mathbf{x}) Q(\mathbf{x})$, and
 - for all $Q^* \in \mathfrak{L}_b^n$ and $\langle j_i \rangle_{i \in [k]} \subseteq [n]$ with $(Q^*(x_0, \dots, x_{n-1}), \psi(x_{j_0}, \dots, x_{j_{k-1}})) \in \mathcal{I}$, we have

$$\mathcal{M} \models (\forall x_0, \dots, x_{n-1}) Q^*(x_0, \dots, x_{n-1}) \rightarrow \psi(x_{j_0}, \dots, x_{j_{k-1}}).$$

Theorem 2.4 (a) and (b) tell us that an invariant completely determines the quantifier-free structure of its arguments. Theorem 2.4 (c) then tells us that each invariant is realized in a structure which is *compatible* with the invariant diagram.

Definition 2.5. We say that an $(\mathfrak{L}_b \cup \mathfrak{L}_c)$ -structure $\mathcal{M}$ is **compatible with $\mathcal{I}$** if it satisfies the following two conditions.

- $\mathcal{M} \models \text{Th}_{\text{SGeo}}(\mathfrak{L}_a, \mathfrak{L}_b, \mathfrak{L}_c)$.
- For $Q \in \mathfrak{L}_b^n$, an atomic or negated atomic $(\mathfrak{L}_b \cup \mathfrak{L}_c)$ -formula $\psi(y_0, \dots, y_{k-1})$, and $\langle j_i \rangle_{i \in [k]} \subseteq [n]$ with $(Q(x_0, \dots, x_{n-1}), \psi(x_{j_0}, \dots, x_{j_{k-1}})) \in \mathcal{I}$, we have

$$\mathcal{M} \models (\forall x_0, \dots, x_{n-1}) Q(x_0, \dots, x_{n-1}) \rightarrow \psi(x_{j_0}, \dots, x_{j_{k-1}}).$$

In other words, $\mathcal{M}$ is compatible with $\mathcal{I}$ if the information contained in any closed tuple in $\mathcal{M}$ agrees with the information the invariant diagram says should hold of any tuple with its invariant.

It will be important, though, that invariants carry information beyond just the quantifier-free information. We therefore need a notion of “information up to complexity β ”.

Definition 2.6. An **invariant system** (on $\mathfrak{L}_b$) is a collection of maps $\boldsymbol{\tau} = \langle \tau_\beta \rangle_{\beta \in \omega_1}$ which satisfy the following.

- (a) For all $\beta \in \omega_1$, we have $\tau_\beta: \mathfrak{L}_b \rightarrow \mathfrak{L}_b$.
- (b) For all $n \geq 1$, for all $\beta \in \omega_1$, we have $\tau_\beta \upharpoonright_{\mathfrak{L}_b^n}: \mathfrak{L}_b^n \rightarrow \mathfrak{L}_b^n$, i.e., for all invariants Q , $\tau_\beta(Q)$ has the same arity as Q .
- (c) For all $\alpha \leq \beta \in \omega_1$, we have $\tau_\alpha \circ \tau_\beta = \tau_\beta \circ \tau_\alpha = \tau_\alpha$.
- (d) If for all $\beta \in \omega_1$ we have $\mathfrak{L}_b[\boldsymbol{\tau}, \beta] = \{Q \in \mathfrak{L}_b : \tau_\beta(Q) = Q\}$, then
 - (i) for all $\beta \in \omega_1$, we have $|\mathfrak{L}_b[\boldsymbol{\tau}, \beta]| \leq \omega$,
 - (ii) for all $\beta \in \omega_1$, we have $\mathfrak{L}_b[\boldsymbol{\tau}, \beta] \neq \mathfrak{L}_b$, and
 - (iii) $\mathfrak{L}_b = \bigcup_{\beta \in \omega_1} \mathfrak{L}_b[\boldsymbol{\tau}, \beta]$.

We also let $\mathfrak{L}_b[\boldsymbol{\tau}, \omega_1] = \mathfrak{L}_b$.

For $\beta \in \omega_1$ we say that the elements of $\mathfrak{L}_b[\boldsymbol{\tau}, \beta]$ are β -**invariants**. We define the $\boldsymbol{\tau}$ -**rank** of an invariant Q , denoted by $\text{rank}_{\boldsymbol{\tau}}(Q)$, to be the least β such that $\tau_\beta(Q) = Q$.

The intuition is that an invariant system gives us a way to take a β -invariant and an $\alpha < \beta$ and return the projection of the β -invariant “down to complexity α ”. Note that all information in $\mathfrak{L}_c$ is preserved under projection, and hence can be thought of as having “complexity 0”. Also note that condition (d) implies $|\mathfrak{L}_b| = \aleph_1$.

Definition 2.7. Suppose $\mathcal{M}$ is an $(\mathfrak{L}_b \cup \mathfrak{L}_c)$ -structure. We define the β -**projection** of $\mathcal{M}$, denoted by $\tau_\beta(\mathcal{M})$, to be the $(\mathfrak{L}_b \cup \mathfrak{L}_c)$ -structure satisfying the following.

- $\mathcal{M}$ and $\tau_\beta(\mathcal{M})$ have the same underlying set M .
- For all relation symbols $R \in \mathfrak{L}_c^n$ and all tuples $a_0, \dots, a_{n-1} \in M$,

$$\tau_\beta(\mathcal{M}) \models R(a_0, \dots, a_{n-1}) \text{ if and only if } \mathcal{M} \models R(a_0, \dots, a_{n-1}).$$

- For all relation symbols $R \in \mathfrak{L}_b^n$ and all tuples $a_0, \dots, a_{n-1} \in M$, we have $\tau_\beta(\mathcal{M}) \models R(a_0, \dots, a_{n-1})$ if and only if there is an $R^* \in \mathfrak{L}_b$ with $\tau_\beta(R^*) = R$ and $\mathcal{M} \models R^*(a_0, \dots, a_{n-1})$.

Suppose $\mathcal{M}$ is an $(\mathfrak{L}_a, \mathfrak{L}_b, \mathfrak{L}_c)$ -structured geometry. We define the τ -**rank** of $\mathcal{M}$, denoted by $\text{rank}_\tau(\mathcal{M})$, to be the least β such that $\tau_\beta(\mathcal{M}) = \mathcal{M}$.

We can think of the β -projection of $\mathcal{M}$ as the structure which is obtained by replacing all invariants in $\mathcal{M}$ with their projections down to β and keeping all other information the same.

It will be important to know when an invariant diagram and an invariant system are compatible.

Definition 2.8. Suppose $\mathcal{I}$ is an invariant diagram and $\tau = \langle \tau_\beta \rangle_{\beta \in \omega_1}$ is an invariant system. We say that $\mathcal{I}$ and τ are **compatible** if for all $\beta \in \omega_1$,

- if $(P(\mathbf{x}), Q(\mathbf{y})) \in \mathcal{I}$ with $P, Q \in \mathfrak{L}_b$ then $(\tau_\beta(P)(\mathbf{x}), \tau_\beta(Q)(\mathbf{y})) \in \mathcal{I}$, and
- if $P \in \mathfrak{L}_b$ and $Q \in \mathfrak{L}_c$ then

$$(P(\mathbf{x}), Q(\mathbf{y})) \in \mathcal{I} \text{ if and only if } (\tau_\beta(P)(\mathbf{x}), Q(\mathbf{y})) \in \mathcal{I}.$$

An invariant diagram and an invariant system are compatible if for all β , taking the β -projection in the invariant system of relations preserves the structure of the invariant diagram.

The following lemma is immediate.

Lemma 2.9. Suppose $\mathcal{I}$ is an invariant diagram and $\tau = \langle \tau_\beta \rangle_{\beta \in \omega_1}$ is an invariant system compatible with $\mathcal{I}$. Then for every $(\mathfrak{L}_b \cup \mathfrak{L}_c)$ -structure $\mathcal{M}$ which is compatible with $\mathcal{I}$ and every $\beta \in \omega_1$, we have that $\tau_\beta(\mathcal{M})$ is compatible with $\mathcal{I}$.

2.3. Density. When studying structures which are compatible with an invariant diagram $\mathcal{I}$, a natural first place to look would be those structures where, if an invariant extension is realized somewhere in the model it must be realized everywhere in the model — specifically on those structures which satisfy the following axioms for all $(Q(\mathbf{xy}), R(\mathbf{x})) \in \mathcal{I}$:

$$(\exists \mathbf{x}^* \mathbf{y}^*) Q(\mathbf{x}^* \mathbf{y}^*) \rightarrow [(\forall \mathbf{x}) R(\mathbf{x}) \rightarrow (\exists \mathbf{y}) Q(\mathbf{xy})].$$

However, this extension property is too strong and would, by a standard back and forth argument, force every model to be ultrahomogeneous with respect to partial isomorphisms between closed sets. Instead we will want a weakened extension property that uses the fact that invariants have projections that lose information.

Definition 2.10. Suppose $\mathcal{I}$ is an invariant diagram and $\tau = \langle \tau_\beta \rangle_{\beta \in \omega_1}$ is an invariant system compatible with $\mathcal{I}$. For $\alpha < \beta \leq \omega_1$ we define the (α, β) -**density theory** of $\mathcal{I}$ and τ , denoted by $\text{Th}_{\text{Den}}(\alpha: \beta, \mathcal{I}, \tau)$, to consist of the following sentences.

- $\text{Th}_{\text{SGeo}}(\mathfrak{L}_a, \mathfrak{L}_b, \mathfrak{L}_c)$.
- Whenever $(Q(x_0, \dots, x_{n-1}), \psi(x_{j_0}, \dots, x_{j_{k-1}})) \in \mathcal{I}$, we have

$$(\forall x_0, \dots, x_{n-1}) Q(x_0, \dots, x_{n-1}) \rightarrow \psi(x_{j_0}, \dots, x_{j_{k-1}}).$$

- For all $Q, R \in \mathfrak{L}_b[\tau, \beta]$ with $(Q(\mathbf{xy}), R(\mathbf{x})) \in \mathcal{I}$, we have

$$(\forall \mathbf{x}) R(\mathbf{x}) \rightarrow (\exists \mathbf{y}) \bigvee_{\substack{Q^* \in \mathfrak{L}_b[\tau, \beta] \\ \tau_\alpha(Q^*) = \tau_\alpha(Q)}} Q^*(\mathbf{xy}).$$

We say that an $(\mathfrak{L}_b \cup \mathfrak{L}_c)$ -structure is (α, β) -**dense** (with respect to $\mathcal{I}$ and τ) if it satisfies the theory $\text{Th}_{\text{Den}}(\alpha: \beta, \mathcal{I}, \tau)$.

When $\beta = \omega_1$ we will omit mention of it.

Intuitively an $(\mathfrak{L}_b \cup \mathfrak{L}_c)$ -structure $\mathcal{M}$ satisfies (α, β) -density with respect to $\mathcal{I}$ and τ if whenever we have a tuple $\mathbf{a} \in \mathcal{M}$ with invariant R that has rank at most β , and there is an invariant Q of rank at most α which not only is an extension of $\tau_\alpha(R)$, but also is the projection of *some* extension of R of rank at most β , then we can find an extension of $\mathbf{a}$ to a closed tuple $\mathbf{ab}$ whose invariant has an α -projection equal to Q .

Note that this is weaker than just saying that whenever $(Q(\mathbf{xy}), R(\mathbf{x})) \in \mathcal{I}$ with both Q and R having rank at most α , then $\tau_\alpha(\mathcal{M})$ satisfies the extension property with respect to $(Q(\mathbf{x}, \mathbf{y}), R(\mathbf{x}))$.

Remark 2.11. Note that the conjunction of $\text{Th}_{\text{Den}}(\alpha, \mathcal{I}, \tau)$ is a sentence in $\mathcal{L}_{\omega_2, \omega}(\mathfrak{L}_b \cup \mathfrak{L}_c)$, not in $\mathcal{L}_{\omega_1, \omega}(\mathfrak{L}_b \cup \mathfrak{L}_c)$, because Theorem 2.10 (c) has an uncountable disjunction. This will not be a problem, however, as our counterexample will not contain the full density axioms themselves but rather will be a “shadow” of a class that satisfies them.

The following lemmas are immediate.

Lemma 2.12. *Suppose $\text{rank}_\tau(\mathcal{M}) \leq \beta$. Then $\mathcal{M} \models \text{Th}_{\text{Den}}(\alpha : \beta, \mathcal{I}, \tau)$ if and only if $\mathcal{M} \models \text{Th}_{\text{Den}}(\alpha, \mathcal{I}, \tau)$.*

Lemma 2.13. *Suppose $\alpha < \beta$ and $\mathcal{M} \models \text{Th}_{\text{Den}}(\beta, \mathcal{I}, \tau)$. Then $\mathcal{M} \models \text{Th}_{\text{Den}}(\alpha, \mathcal{I}, \tau)$.*

Lemma 2.14. *Suppose $\alpha \leq \beta$, $\mathcal{M}$ is compatible with $\mathcal{I}$ and $\tau_\beta(\mathcal{M}) \models \text{Th}_{\text{Den}}(\alpha, \mathcal{I}, \tau)$. Then $\mathcal{M} \models \text{Th}_{\text{Den}}(\alpha, \mathcal{I}, \tau)$.*

2.4. Back and Forth Systems. We now introduce one of the fundamental tools for studying structured geometries in a theory.

Definition 2.15. Suppose $\mathcal{M}$ and $\mathcal{N}$ are $\mathfrak{L}_c$ -structures which have an $\mathfrak{L}_a$ -geometry. For $\alpha \in \text{ORD}$, a **back and forth system** (of length α) between $\mathcal{M}$ and $\mathcal{N}$ is a sequence $\langle I_i \rangle_{i \leq \alpha}$ which satisfies the following.

- (a) For $i < j \leq \alpha$ we have $I_j \subseteq I_i$ and $I_\alpha \neq \emptyset$.
- (b) For each $i \leq \alpha$, the set I_i consists of pairs of closed finite tuples $(\mathbf{a}, \mathbf{b})$ of the same length, where $\mathbf{a} \subseteq \mathcal{M}$ and $\mathbf{b} \subseteq \mathcal{N}$.
- (c) If $\ell \leq \alpha$ and $(\langle a_i \rangle_{i \in [n]}, \langle b_i \rangle_{i \in [n]}) \in I_\ell$, then for all relation symbols $R \in \mathfrak{L}_c^k$ and all sequences $\langle j_i \rangle_{i \in [k]} \subseteq [n]$, we have

$$\mathcal{M} \models R(a_{j_0}, \dots, a_{j_{k-1}}) \text{ if and only if } \mathcal{N} \models R(b_{j_0}, \dots, b_{j_{k-1}}).$$

- (d) For each $i < j \leq \alpha$ and each pair $(\mathbf{a}, \mathbf{b}) \in I_j$,
 - (i) if $\mathbf{a}^* \subseteq \mathcal{M}$ and $\mathbf{aa}^*$ is closed then there is a $\mathbf{b}^* \subseteq \mathcal{N}$ such that $(\mathbf{aa}^*, \mathbf{bb}^*) \in I_i$, and
 - (ii) if $\mathbf{b}^* \subseteq \mathcal{N}$ and $\mathbf{bb}^*$ is closed then there is a $\mathbf{a}^* \subseteq \mathcal{M}$ such that $(\mathbf{aa}^*, \mathbf{bb}^*) \in I_i$.

We say that $\mathcal{M}$ and $\mathcal{N}$ are **α -equivalent**, denoted by $\mathcal{M} \equiv_\alpha \mathcal{N}$, if there is a back and forth system of length α , each component of which is non-empty.

If $\langle I_i \rangle_{i \leq \beta}$ is a back and forth system between $\mathcal{M}$ and $\mathcal{N}$ and $(\mathbf{a}, \mathbf{b}) \in I_\alpha$ (for $\alpha \leq \beta$) then we can think of $\mathbf{a}$ and $\mathbf{b}$ as being “indistinguishable up to complexity α ”. Note, though, that the relation of $\equiv_\beta$ need not be an equivalence relation (although it is symmetric). One way to fix this would be to require that $(\langle \rangle, \langle \rangle) \in I_i$ for each $i \leq \beta$, and indeed this is what is done in classic back and forth arguments. But as in our context the empty tuple is not closed, this would remove the condition that all tuples in our back and forth systems are closed.

The following is a generalization of a standard result.

Proposition 2.16. *Suppose Φ is a class of geometric $\mathfrak{L}_c$ -structures closed under isomorphism such that for all $\beta \in \omega_1$, all $\equiv_\beta$ -antichains in Φ are countable. Then there does not exist a perfect set of non-isomorphic models in Φ .*

Proof. Suppose to get a contradiction that there is a perfect set P of non-isomorphic models in Φ . Let BO be the collection of codes for countable linear orders $(L, \leq)$ such that there are distinct $\mathcal{M}, \mathcal{N} \in P$ and a sequence $\langle I_i \rangle_{i \in L}$ which satisfy the conditions of Theorem 2.15, each element of which

is non-empty. If L is ill-founded then by a standard back and forth argument we must have $\mathcal{M} \cong \mathcal{N}$. Therefore, because no two elements of P are isomorphic, any linear order in BO must be well-founded.

But BO is a Σ_1^1 -class as the conditions in Theorem 2.15 are Borel in an appropriate space. So by Σ_1^1 -bounding [Jec03, Corollary 25.14] there must be a $\gamma < \omega_1$ such that all orders in BO have order type at most γ . Therefore if $\mathcal{M}, \mathcal{N} \in P$ and $\mathcal{M} \equiv_{\gamma+1} \mathcal{N}$ we must have $\mathcal{M} = \mathcal{N}$. Hence P is an antichain in $\equiv_{\gamma+1}$. But $|P| = 2^\omega$, getting us our contradiction. $\square$

2.5. Weak Knight Systems. We now introduce one of the basic structures in this section. The intuition is that a weak Knight system consists, for each β , of a collection of models of rank at most β which satisfy density below β and where being in the weak Knight system is preserved by projections in the invariant system.

Definition 2.17. Suppose $\mathcal{I}$ is an invariant diagram and τ is an invariant system compatible with $\mathcal{I}$. A **weak Knight system** (for $\mathcal{I}$ and τ) is a sequence $\langle K_\beta \rangle_{\beta \in \omega_1}$ satisfying the following for each $\beta \in \omega_1$.

- Each K_β is an isomorphism closed collection of countable $(\mathcal{L}_b \cup \mathcal{L}_c)$ -structures compatible with $\mathcal{I}$ and with rank at most β .
- For all $\alpha < \beta$ and all $\mathcal{M} \in K_\beta$, $\mathcal{M} \models \text{Th}_{\text{Den}}(\alpha, \mathcal{I}, \tau)$, i.e. $\mathcal{M}$ satisfies α -density.
- For all $\alpha < \beta$ and all $\mathcal{M} \in K_\beta$, $\tau_\alpha(\mathcal{M}) \in K_\alpha$.

We will let $K = \bigcup_{\beta \in \omega_1} K_\beta$. We call $K_{-\infty} = \{\mathcal{M} \upharpoonright_{\mathcal{L}_c} : \mathcal{M} \in K\}$ the **base** of K .

Note that for a weak Knight system we do not assume for each β that K_β contains at most countably many non-isomorphic structures.

One of the main advantages of weak Knight systems is that they give a way to produce back and forth systems between their models.

Proposition 2.18. *Suppose that*

- $\langle K_\beta \rangle_{\beta \in \omega_1}$ is a weak Knight system (for $\mathcal{I}$ and τ),
- $\mathcal{M}, \mathcal{N} \in K_\beta$, and
- there is an invariant $Q \in \mathcal{L}_b$ such that $\mathcal{M} \models (\exists x) Q(x)$ and $\mathcal{N} \models (\exists x) Q(x)$.

Then $\mathcal{M} \upharpoonright_{\mathcal{L}_c} \equiv_\beta \mathcal{N} \upharpoonright_{\mathcal{L}_c}$.

Proof. For $i \leq \beta$ let I_i be the collection of pairs $(\mathbf{a}, \mathbf{b})$ such that

- $\mathbf{a} \subseteq \mathcal{M}$ and $\mathbf{b} \subseteq \mathcal{N}$ are closed and of the same length, and
- there is a $Q \in \mathcal{L}_b$ such that $\tau_i(\mathcal{M}) \models Q(\mathbf{a})$ and $\tau_i(\mathcal{N}) \models Q(\mathbf{b})$.

It suffices to show that $\langle I_i \rangle_{i \leq \beta}$ is a back and forth system. The fact that when $i < j \leq \beta$ we have $I_j \subseteq I_i$ follows from the compatibility of τ and $\mathcal{I}$ and the fact that $\tau_i = \tau_i \circ \tau_j$. Further, because $\mathcal{M}, \mathcal{N} \in K_\beta$ we have $\max\{\text{rank}_\tau(\mathcal{M}), \text{rank}_\tau(\mathcal{N})\} \leq \beta$ and so $\tau_\beta(\mathcal{M}) = \mathcal{M}$ and $\tau_\beta(\mathcal{N}) = \mathcal{N}$. So if $\mathcal{M} \models Q(\mathbf{a})$ and $\mathcal{N} \models Q(\mathbf{b})$ we have $(\mathbf{a}, \mathbf{b}) \in I_\beta$. Hence Theorem 2.15 (a) holds.

Theorem 2.15 (b) is immediate and Theorem 2.15 (c) follows from the fact that invariants determine the quantifier-free $\mathcal{L}_c$ -structure on their arguments.

We now show Theorem 2.15 (d.i); the proof of Theorem 2.15 (d.ii) is identical. Suppose $i < j \leq \beta$, $(\mathbf{a}, \mathbf{b}) \in I_j$ and $\mathbf{a}\mathbf{a}^*$ is closed in $\mathcal{M}$. Let $P_{\mathbf{a}}, P_{\mathbf{b}}, Q_{\mathbf{a}\mathbf{a}^*} \in \mathcal{L}_b$ be such that $\tau_j(\mathcal{M}) \models P_{\mathbf{a}}(\mathbf{a})$, $\tau_j(\mathcal{M}) \models Q_{\mathbf{a}\mathbf{a}^*}(\mathbf{a}\mathbf{a}^*)$ and $\tau_j(\mathcal{N}) \models P_{\mathbf{b}}(\mathbf{b})$. We then have $(Q_{\mathbf{a}\mathbf{a}^*}, P_{\mathbf{a}}) \in \mathcal{I}$ and $P_{\mathbf{a}} = P_{\mathbf{b}}$.

Because $\tau_j(\mathcal{N}) \in K_j$ we have that $\tau_j(\mathcal{N})$ satisfies i -density. Therefore there must be some $\mathbf{b}^* \subseteq \mathcal{N}$ and $Q_{\mathbf{b}\mathbf{b}^*} \in \mathcal{L}_b$ such that $\tau_j(\mathcal{N}) \models Q_{\mathbf{b}\mathbf{b}^*}(\mathbf{b}\mathbf{b}^*)$ and $\tau_i(Q_{\mathbf{b}\mathbf{b}^*}) = \tau_i(Q_{\mathbf{a}\mathbf{a}^*})$. But then $(\mathbf{a}\mathbf{a}^*, \mathbf{b}\mathbf{b}^*) \in I_i$. As $(\mathbf{a}, \mathbf{b})$ was arbitrary this implies Theorem 2.15 (d.i) and so $\langle I_i \rangle_{i \leq \beta}$ is a back and forth system. $\square$

In most situations we will have that “most” structures in the base of a weak Knight system will have lifts of high rank. In this situation the existence of a weak Knight system ensures that the base does not contain a perfect set of non-isomorphic models.

Theorem 2.19. *Suppose that*

- $K = \langle K_\beta \rangle_{\beta \in \omega_1}$ is a weak Knight system, and

(b) for all $\beta \in \omega_1$,

$$\{\mathcal{M} \in K_{-\infty} : (\forall \mathcal{N} \in K_\beta) \mathcal{N} \upharpoonright_{\mathfrak{L}_c} \neq \mathcal{M}\}$$

has countably many isomorphism classes.

Then $K_{-\infty}$ does not contain a perfect set of non-isomorphic models.

Proof. Fix $\beta \in \omega_1$. Let

$$A_\beta = \{\mathcal{M} \in K_{-\infty} : (\forall \mathcal{N} \in K_\beta) \mathcal{N} \upharpoonright_{\mathfrak{L}_c} \neq \mathcal{M}\}.$$

By (b) A_β has only countably many isomorphism classes.

For $\mathcal{M}, \mathcal{N} \in K_\beta$ define $\mathcal{M} \sim \mathcal{N}$ if there is some $Q \in \mathfrak{L}_b$ such that $\mathcal{M} \models (\exists \mathbf{x}) Q(\mathbf{x})$ and $\mathcal{N} \models (\exists \mathbf{x}) Q(\mathbf{x})$. Note that every relation symbol in $\mathfrak{L}_b$ which is realized in a structure in K_β must have rank at most β and hence be in $\mathfrak{L}_b[\tau, \beta]$. But $\mathfrak{L}_b[\tau, \beta]$ is countable and so any $\sim$ -antichain must be countable.

By the definition of A_β , any model $\mathcal{M} \in K_{-\infty} \setminus A_\beta$ must have a lift $\mathcal{M}^* \in K_\beta$. By Theorem 2.18 if B is a $\equiv_\beta$ -antichain among elements of $K_{-\infty} \setminus A_\beta$ then there must be an $\sim$ -antichain B^+ such that $B = \{\mathcal{M} \upharpoonright_{\mathfrak{L}_c} : \mathcal{M} \in B^+\}$. Therefore every $\equiv_\beta$ -antichain in $K_{-\infty} \setminus A_\beta$ must be countable. Hence, because A_β is countable up to isomorphism, any $\equiv_\beta$ -antichain in $K_{-\infty}$ must be countable as well. But then by Theorem 2.16, $K_{-\infty}$ does not contain a perfect set of non-isomorphic models. $\square$

2.6. Knight Systems and Knight Classes. We now introduce the central object in this section, the Knight system. This is a weak Knight system with extra properties which ensure that its base satisfies the conditions of being a counterexample to Vaught's conjecture (provided the base is the class of countable models of a sentence of $\mathcal{L}_{\omega_1, \omega}$). First, though, we need a few related definitions.

Definition 2.20. Suppose $\mathbf{K} = \langle K_\beta \rangle_{\beta \in \omega_1}$ is a weak Knight system (for $(\mathcal{I}, \tau)$) and $\beta \in \omega_1$. We say that $\mathcal{M} \in K_\beta$ is **maximal at level β** (for $\mathbf{K}$) if whenever $\beta^* > \beta$ and $\mathcal{M}^* \in K_{\beta^*}$ and $\tau_\beta(\mathcal{M}^*) = \mathcal{M}$, then $\mathcal{M} = \mathcal{M}^*$. We say that $\mathcal{M}$ is **maximal** if it is maximal at some level.

We say $\mathbf{K}$ has **few maximal models** if for all $\beta \in \omega_1$, up to isomorphism, there are countably many models which are maximal at β .

Definition 2.21. Suppose that $\mathbf{K} = \langle K_\beta \rangle_{\beta \in \omega_1}$ is a weak Knight system (for $(\mathcal{I}, \tau)$), that $\mathcal{M}$ is an $\mathfrak{L}_c$ -structure, and that $\mathbf{a} \subseteq \mathcal{M}$ is closed. We say that $\mathbf{a}$ is **eventually determined in $\mathcal{M}$ with respect to $\mathbf{K}$** if there are $\alpha_{\mathbf{a}} \in \omega_1$ and $Q \in \mathfrak{L}_b$ satisfying the following conditions.

- There is some lift $\mathcal{M}^*$ of $\mathcal{M}$ with $\mathcal{M}^* \in K_{\alpha_{\mathbf{a}}}$ and $\mathcal{M}^* \models Q(\mathbf{a})$.
- Whenever $\alpha^* > \alpha_{\mathbf{a}}$ and $\mathcal{M}^*$ is a lift of $\mathcal{M}$ with $\mathcal{M}^* \in K_{\alpha^*}$ then $\mathcal{M}^* \models Q(\mathbf{a})$.

When the structure and weak Knight system are understood, we simply say that a tuple is “eventually determined”.

We say that $\mathcal{M}$ is **eventually determined with respect to $\mathbf{K}$** if every closed tuple in $\mathcal{M}$ is eventually determined.

In other words, a tuple is eventually determined if all lifts past a certain level have to agree on the invariant of the tuple, and a structure is eventually determined if all closed tuples are eventually determined. Note that a structure being eventually determined does not necessarily imply that there is a maximal lift. This is because the levels at which the tuples stabilize could have a supremum which is a limit ordinal, and there might be no lift of the structure at that limit level.

Lemma 2.22. *If $\mathcal{M}$ is eventually determined then there is an $\alpha_{\mathcal{M}} < \omega_1$ such that whenever $\mathcal{M}^* \in K$ is a lift of $\mathcal{M}$ we have $\text{rank}_\tau(\mathcal{M}^*) \leq \alpha_{\mathcal{M}}$.*

Proof. Let $\langle \mathbf{a}_i \rangle_{i \in \omega}$ be an enumeration of all closed tuples of $\mathcal{M}$. For $i \in \omega$ let $\alpha_{\mathbf{a}_i} \in \omega_1$ be minimal such that it satisfies the conditions in Theorem 2.21. Note that as K_β is closed under isomorphism for each β , the value $\alpha_{\mathbf{a}_i}$ depends only on the isomorphism type of $(\mathcal{M}, \mathbf{a}_i)$. Let $\alpha_{\mathcal{M}} = \sup\{\alpha_{\mathbf{a}_i} + 1 : i \in \omega\}$.

Suppose to get a contradiction there is a $\mathcal{M}^* \in K$ which is a lift of $\mathcal{M}$ with $\text{rank}_\tau(\mathcal{M}^*) > \alpha_{\mathcal{M}}$. Then for each $i \in \omega$ and each $Q \in \mathfrak{L}_b$, since $\mathbf{a}_i$ is eventually determined, we have that $\mathcal{M}^* \models Q(\mathbf{a}_i)$ if and only if $\tau_{\alpha_{\mathcal{M}}}(\mathcal{M}^*) \models Q(\mathbf{a}_i)$. Therefore $\mathcal{M}^* = \tau_{\alpha_{\mathcal{M}}}(\mathcal{M}^*)$ and so $\text{rank}_\tau(\mathcal{M}^*) \leq \alpha_{\mathcal{M}}$, yielding a contradiction. $\square$

Definition 2.23. Suppose $\mathbf{K} = \langle K_\beta \rangle_{\beta \in \omega_1}$ is a weak Knight system (for $\mathcal{I}$ and τ). We say $\mathbf{K}$ is a **Knight system** if

- (a) $K_{-\infty} = \{\mathcal{M} \upharpoonright_{\mathcal{L}_c} : \mathcal{M} \in K \text{ and } \mathcal{M} \text{ is maximal with respect to } \mathbf{K}\}$,
- (b) $\mathbf{K}$ has few maximal models,
- (c) if $\mathcal{M} \in K_{-\infty}$ then $\mathcal{M}$ is eventually determined with respect to $\mathbf{K}$, and
- (d) for every invariant $Q \in \mathcal{L}_b$ there is some $\mathcal{M} \in K$ where $\mathcal{M} \models (\exists \mathbf{x}) Q(\mathbf{x})$.

Definition 2.24. We say a collection Φ of $\mathcal{L}_c$ -structures is a **Knight class** if there is an invariant diagram $\mathcal{I}$, a compatible invariant system $\tau = \langle \tau_\beta \rangle_{\beta \in \omega_1}$, and a Knight system $\mathbf{K} = \langle K_\beta \rangle_{\beta \in \omega_1}$ (for $(\mathcal{I}, \tau)$) such that $\Phi = K_{-\infty}$. We say that $\mathbf{K}$ **witnesses** that Φ is a Knight class.

We then have the following theorem, which is the culmination of all the results in this section.

Theorem 2.25. *Suppose Φ is a Knight class.*

- (a) Φ has precisely $\aleph_1$ many countable models up to isomorphism.
- (b) There does not exist a perfect set of non-isomorphic models in Φ .

Proof. As $\mathbf{K}$ has few maximal models, there are at most $\aleph_1$ many maximal models for $\mathbf{K}$ up to isomorphism. But by Theorem 2.23 (a) every element of Φ is the reduct of (at least) one maximal model for $\mathbf{K}$. Therefore Φ has at most $\aleph_1$ many models up to isomorphism.

For all $\mathcal{M} \in \Phi$, let $\alpha_{\mathcal{M}}$ be as in Theorem 2.22. If $Q \in \mathcal{L}_b$ and $\text{rank}_\tau(Q) > \alpha_{\mathcal{M}}$, then Q is not realized in any lift of $\mathcal{M}$. Therefore, as there are Q with arbitrarily large ranks, $\{\alpha_{\mathcal{M}} : \mathcal{M} \in \Phi\}$ must be cofinal in ω_1 and hence have size $\aleph_1$. Therefore Φ must have at least $\aleph_1$ many models up to isomorphism and so (a) holds.

Fix $\beta \in \omega_1$. Let

$$A_\beta = \{\mathcal{M} \in \Phi : \text{no lift of } \mathcal{M} \text{ is maximal at any } \delta \geq \beta\}.$$

Because $\mathbf{K}$ has few maximal models, A_β has only countably many isomorphism classes. But then $\{\mathcal{M} \in \Phi : (\forall \mathcal{N} \in K_\beta) \mathcal{N} \upharpoonright_{\mathcal{L}_c} \neq \mathcal{M}\}$ also has countably many models up to isomorphism. Therefore, by Theorem 2.19, Φ does not have a perfect set of non-isomorphic models and (b) holds. $\square$

Definition 2.26. We say that a sentence $\sigma \in \mathcal{L}_{\omega_1, \omega}(\mathcal{L}_c)$ is a **Knight sentence** if the collection of countable models of σ forms a Knight class.

The following is immediate from Theorem 2.25.

Theorem 2.27. *If $\sigma \in \mathcal{L}_{\omega_1, \omega}(\mathcal{L}_c)$ is a Knight sentence then σ has $\aleph_1$ many countable models up to isomorphism and does not have a perfect set of non-isomorphic models. Hence σ is a counterexample to Vaught's conjecture for $\mathcal{L}_{\omega_1, \omega}$.*

3. TOY EXAMPLE

We will now introduce a toy example. This example will produce a Knight class, even though it will not produce a Knight sentence. Walking through how we create the Knight system, and seeing what fails, will hopefully help provide insight into how the eventual counterexample behaves.

3.1. Labeled Geometric Trees. We now introduce the structures of our example. They will be trees with labels that are meant to capture information about the rank of elements.

Definition 3.1. Let $\mathcal{L}_a = \{P_n\}_{n \geq 1}$ where, for each $n \geq 1$, the relation symbol P_n is of arity n . Let $\mathcal{L}_c = \mathcal{L}_a \cup \{\leq\}$ where $\leq$ is a binary relation symbol.

We say an $\mathcal{L}_c$ -structure $\mathcal{M}$ is a **geometric tree** if it satisfies the following.

- $\leq$ is a partial order with a greatest element called the root.
- For all $a \in \mathcal{M}$, the set $\{a' \in \mathcal{M} : a' \geq a\}$ is finite and linearly ordered.
- For $n \geq 1$, we have

$$(\forall x_0, \dots, x_{n-1}) P_n(x_0, \dots, x_{n-1}) \leftrightarrow \left[\left(\bigwedge_{i < j \in [n]} x_i \neq x_j \right) \wedge (\forall y) \left[\left(\bigvee_{i \in [n]} y \geq x_i \right) \rightarrow \left(\bigvee_{i \in [n]} y = x_i \right) \right] \right].$$

If $\mathcal{M}$ is a geometric tree and $a, b \in \mathcal{M}$, we will say $\mathcal{M} \models a <_1 b$ if $\mathcal{M} \models a < b$ and $\mathcal{M} \models (\forall c) a \leq c \leq b \rightarrow (c = a \vee c = b)$, i.e., a is an immediate predecessor of b .

In other words, a geometric tree is a tree where the closed sets are finite upwards closed sets. Note that in a geometric tree, the tree ordering is infinitarily definable from the geometry.

The following is immediate, and says that geometric trees have an $\mathfrak{L}_a$ -geometry.

Lemma 3.2. *If $\mathcal{M}$ is a geometric tree then $\mathcal{M} \models \text{Th}_{\text{Geo}}(\mathfrak{L}_a)$.*

The following is an important function on geometric trees.

Definition 3.3. Suppose $\mathcal{M}$ is a countable geometric tree. The $\mathcal{M}$ -**rank** is a map $\text{rank}_{\mathcal{M}}: M \rightarrow \omega_1 \cup \{\infty\}$ where $\text{rank}_{\mathcal{M}}(a)$ is defined as follows.

- If $\mathcal{M} \models (\forall b) b \not<_1 a$ then $\text{rank}_{\mathcal{M}}(a) = 0$.
- If there is a sequence $\langle a_i \rangle_{i \in \omega}$ such that for all $i \in \omega$ we have $\mathcal{M} \models a_{i+1} <_1 a_i$ and $a_0 = a$, then $\text{rank}_{\mathcal{M}}(a) = \infty$.
- Otherwise, $\text{rank}_{\mathcal{M}}(a) = \sup\{\text{rank}_{\mathcal{M}}(b) + 1 : b <_1 a\}$.

Definition 3.4. By a **labeled tree** we mean a pair $\langle (T, \leq), \mu \rangle$ where

- $(T, \leq)$ is a rooted tree,
- $\mu: T \rightarrow \omega_1 \cup \{\infty\}$,
- for all $a < b \in T$, we have $\mu(a) < \mu(b)$, where $\infty < \infty$.

We call T the **underlying set** of the labeled tree and we call μ the **labeling**. We say that a labeled tree is **finite** when its underlying set is.

If $T_0 \subseteq T$ is an upwards closed set and $\mu_0 = \mu \upharpoonright_{T_0}$, then we say that $\langle (T, \leq), \mu \rangle$ **extends** $\langle (T_0, \leq \upharpoonright_{T_0 \times T_0}), \mu_0 \rangle$.

It is useful to think of a labeled tree as assigning a value $\alpha \in \omega_1$ when we know the rank of the node is α , but assigning the value ∞ if either the node is ill-founded or we don't yet know what the actual rank is. Our invariants will be built out of these finite labeled trees.

Definition 3.5. For $n \geq 1$ let $\mathfrak{L}_b^n = \{Q_\zeta : \zeta \text{ is a finite labeled tree with underlying set } [n]\}$.

We will be interested in an invariant diagram which will allow all “local labelings” to be combined into a “global labeling”.

Definition 3.6. Let $\mathcal{I}_{\text{tr}}$ be the invariant diagram such that whenever $\zeta = \langle ([n], \leq), \mu \rangle$ is a finite labeled tree and $P \in \mathfrak{L}_b \cup \mathfrak{L}_c$ then $(Q_\zeta(x_0, \dots, x_{n-1}), P(x_{i_0}, \dots, x_{i_{k-1}})) \in \mathcal{I}_{\text{tr}}$ precisely when one of the following holds.

- P is $\leq$ and $i_0 \leq i_1$,
- P is P_k , $i_0, \dots, i_{k-1}$ are distinct elements and $\{i_j\}_{j \in [k]}$ is an upwards closed set in $([n], \leq)$,
- P is Q_{ζ^*} where $\zeta^* = \langle ([k], \leq_{\zeta^*}), \nu \rangle$, the map $i_\ell \mapsto \ell$ is an isomorphism between $([n], \leq_\zeta) \upharpoonright_{\{i_j\}_{j \in [k]}}$ and $([k], \leq_{\zeta^*})$, the sequence $\{i_j\}_{j \in [k]}$ is upwards closed, and $\nu(j) = \mu(i_j)$ for all $j \in [k]$.

We call a structure compatible with $\mathcal{I}_{\text{tr}}$ a **labeled geometric tree**.

Definition 3.7. Suppose $\mathcal{M}$ is an $(\mathfrak{L}_b \cup \mathfrak{L}_c)$ -structure. Let $\text{glab}_{\mathcal{M}}$ be the collection of pairs $(a, \alpha) \in M \times (\omega_1 \cup \{\infty\})$ such that there is a $\zeta = \langle ([n], \leq), \mu \rangle$ and a sequence $a_0, \dots, a_{n-1} \in M$ such that

- $\mathcal{M} \models Q_\zeta(a_0, \dots, a_{n-1})$,
- $a = a_k$ for some $k \in [n]$, and
- $\mu(k) = \alpha$.

Because $\mathcal{I}_{\text{tr}}$ forces relations Q_ζ and Q_{ζ^*} to “assign” the same label to any common argument, we have the following immediate lemma.

Lemma 3.8. *Suppose $\mathcal{M}$ is compatible with $\mathcal{I}_{\text{tr}}$. Then $\mathcal{M} \upharpoonright_{\mathfrak{L}_c}$ is a geometric tree, $\text{glab}_{\mathcal{M}}$ is the graph of a function $\text{lab}_{\mathcal{M}}$ such that $\langle (M, \leq), \text{lab}_{\mathcal{M}} \rangle$ is a labeled tree, and $\text{lab}_{\mathcal{M}}(a) \geq \text{rank}_{\mathcal{M}}(a)$ for all $a \in M$.*

3.2. Knight System: Attempt 1. We now attempt to build a Knight system over a class of trees. As we will see, the obvious thing to try will fail. However, seeing how it fails will be helpful for understanding the eventual counterexample. In particular, as the next example shows, if we intend for invariants to capture the relevant information about their arguments, we will have to restrict attention to some special collection of trees.

Example 3.9. Let $A \subseteq \omega$. Let $\mathcal{M}_A$ be the structured tree satisfying the following.

- $M_A = \{\langle r \rangle\} \cup \{\langle r, i \rangle : i \in \omega\} \cup \{\langle r, i, j \rangle : i \in A \text{ and } j \in [i+2]\} \cup \{\langle r, i, 0 \rangle : i \notin A\}$.
- $a \leq b$ if b is an initial segment of a .
- $\text{lab}_{\mathcal{M}_A} = \text{rank}_{\mathcal{M}_A}$.

All elements of the form $\langle r, i \rangle$ have rank 1 and so all tuples of the form $\langle \langle r \rangle, \langle r, i \rangle \rangle$ satisfy the same invariant. However, we can only extend a map from $\{\langle r \rangle, \langle r, i \rangle\}$ to $\{\langle r \rangle, \langle r, j \rangle\}$ to an automorphism of the structure if either $i = j$ or $\{i, j\} \cap A = \emptyset$. Further, if $A, B \subseteq \omega$ are distinct, then $\mathcal{M}_A \not\cong \mathcal{M}_B$. As such, these trees contain a great deal more information than is captured by their labels.

Definition 3.10. Suppose $\beta \in \omega_1$. We say a countable geometric tree $\mathcal{M}$ is β -full if for all $\gamma < \beta$,

$$(\forall a \in M) \text{rank}_{\mathcal{M}}(a) > \gamma \rightarrow (\exists^\omega b) b <_1 a \text{ and } \text{rank}_{\mathcal{M}}(b) = \gamma.$$

We say $\mathcal{M}$ is **full** if it is $\theta_{\mathcal{M}}$ -full where $\theta_{\mathcal{M}} = \inf\{\beta : \text{range}(\text{rank}_{\mathcal{M}}) \subseteq \beta \cup \{\infty\}\}$, and every node of rank ∞ has infinitely many immediate predecessors of rank ∞ . Let Φ be the collection of full geometric trees.

The following is then straightforward.

Lemma 3.11. *Suppose that*

- $\mathcal{M}, \mathcal{N}$ are full trees,
- $\text{lab}_{\mathcal{M}} = \text{rank}_{\mathcal{M}}$ and $\text{lab}_{\mathcal{N}} = \text{rank}_{\mathcal{N}}$,
- $\text{range}(\text{rank}_{\mathcal{M}}) \subseteq \text{range}(\text{rank}_{\mathcal{N}})$,
- $A \subseteq M$ is a finite closed set, and
- $i: A \rightarrow \mathcal{N}$ is an embedding of labeled geometric trees.

Then i extends to an embedding of labeled geometric trees from $\mathcal{M}$ to $\mathcal{N}$.

Theorem 3.11 tells us that among full trees the invariant of the tuple completely determines the structure of a tuple. For this reason we will want to focus our attention on full trees.

The fact that trees compatible with $\mathcal{I}_{\text{tr}}$ always have a global labeling function suggests a natural collection of projection functions, i.e. functions τ_{β}^- which replace all rank values of β and above with ∞ and keep the rest the same.

Definition 3.12. We define the β -truncation function, denoted by tr_{β} , to be the function $\text{tr}_{\beta}: \omega_1 \cup \{\infty\} \rightarrow \omega_1 \cup \{\infty\}$ where

$$\text{tr}_{\beta}(\alpha) = \begin{cases} \alpha & \text{if } \alpha < \beta, \\ \infty & \text{otherwise.} \end{cases}$$

We then define the following invariant system.

Definition 3.13. For $\beta \in \omega_1$ let $\tau_{\beta}^-: \mathfrak{L}_b \rightarrow \mathfrak{L}_b$ be the map given by $\tau_{\beta}^-(Q_{\zeta}) = Q_{\zeta^*}$ for $\zeta = \langle \langle T, \leq \rangle, \mu \rangle$, where $\zeta^* = \langle \langle T, \leq \rangle, \text{tr}_{\beta} \circ \mu \rangle$.

The following lemma then holds, because whenever $\alpha < \beta$, we have $\text{tr}_{\alpha} \circ \text{tr}_{\beta} = \text{tr}_{\beta} \circ \text{tr}_{\alpha} = \text{tr}_{\alpha}$.

Lemma 3.14. $\tau^- = \langle \tau_{\beta}^- \rangle_{\beta \in \omega_1}$ is an invariant system compatible with $\mathcal{I}_{\text{tr}}$.

We would now like to define a Knight system with respect to $\mathcal{I}_{\text{tr}}$ and τ^- . Unfortunately, as we will see, no such Knight system can exist, and we will eventually need to modify the definition of τ^- . But it will be instructive to see *why* such a Knight system can't exist, before we show how to create one for a closely related invariant system.

First notice that if $\mathbf{K}^- = \langle K_\beta^- \rangle_{\beta \in \omega_1}$ is a weak Knight system for τ^- then for any $\mathcal{M} \in K_\beta^-$ we must have (a) $\tau_\beta^-(\mathcal{M}) = \mathcal{M}$ and (b) for all $\alpha < \beta$, the structure $\mathcal{M}$ is α -dense (with respect to $\mathcal{I}_{\text{tr}}$ and τ^-).

Definition 3.15. Suppose $\beta \in \omega_1$. We say that a labeled geometric tree $\mathcal{M}$ has **labels cut off at β** if

$$(\forall a \in M) \text{lab}_{\mathcal{M}}(a) \geq \beta \rightarrow \text{lab}_{\mathcal{M}}(a) = \infty.$$

The following is immediate.

Lemma 3.16. Suppose $\mathcal{M}$ is a labeled geometric tree. Then the following are equivalent.

- $\mathcal{M}$ has labels cut off at β .
- $\tau_\beta^-(\mathcal{M}) = \mathcal{M}$.

Therefore any $\mathcal{M} \in K_\beta^-$ must have labels cut off at β .

Definition 3.17. Suppose $\beta \leq \omega_1$. A labeled geometric tree $\mathcal{M}$ has **β -true labels** if, for all $a \in M$, we have

$$\text{rank}_{\mathcal{M}}(a) < \beta \rightarrow \text{rank}_{\mathcal{M}}(a) = \text{lab}_{\mathcal{M}}(a).$$

We say that $\mathcal{M}$ has **true labels** if it has ω_1 -true labels.

We then have the following characterization of when a labeled geometric tree is β -dense (with respect to $\mathcal{I}_{\text{tr}}$ and τ^-).

Proposition 3.18. Suppose $\mathcal{M}$ is a countable labeled geometric tree compatible with $\mathcal{I}_{\text{tr}}$. Then the following are equivalent.

- (a) $\mathcal{M}$ is β -dense (with respect to $\mathcal{I}_{\text{tr}}$ and τ^-),
 - (b) $\mathcal{M} \upharpoonright_{\mathcal{L}_c}$ is β -full, $\mathcal{M}$ has β -true labels and for all $n \in \omega$ and $a \in M$,
- ($\mathfrak{A}$) $\beta + n < \text{lab}_{\mathcal{M}}(a) \rightarrow |\{b \in M : b <_1 a \text{ and } \text{lab}_{\mathcal{M}}(b) \geq \beta + n\}| = \omega.$

Proof. Assume (a). Note if $\tau_\beta^-(Q_\zeta) = Q_{\zeta^*}$ then the labeling of ζ^* must have range in $\beta \cup \{\infty\}$. That $\mathcal{M}$ has β -true labels then follows from an easy induction. The fact that $\mathcal{M}$ is β -full follows from the fact that the collection of invariants with a labeling whose range is contained in $\beta \cup \{\infty\}$ has amalgamation, in the obvious sense. By density we may add arbitrarily many immediate predecessors labeled $\gamma < \beta$ whenever γ is less than the label of a node. Further these labels are preserved by projection.

Now suppose $a \in M$, $n \in \omega$ and $\beta + n < \text{lab}_{\mathcal{M}}(a)$. Let $\langle a_i \rangle_{i \in [k]}$ be such that $a = a_{k-1} <_1 a_{k-2} <_1 \dots <_1 a_0$ where a_0 is the root. Let ζ be such that $\mathcal{M} \models Q_\zeta(a_0, \dots, a_{k-1})$. Let $\zeta^- = \langle ([n+k+1], \leq), \mu^- \rangle$ where $\mu^-(k+i) = \beta + (n-i)$ for $0 \leq i \leq n$ and $k+n <_1^- k+n-1 <_1^- \dots <_1^- k <_1^- k-1$. Because $\tau_\beta^-(Q_{\zeta^-})(x_0, \dots, x_{k+n})$ implies the label of each of $x_k, \dots, x_{k+n}$ is ∞ , if ζ^* is such that $\tau_\beta^-(Q_{\zeta^-}) = \tau_\beta^-(Q_{\zeta^*})$ we have $Q_{\zeta^*}(x_0, \dots, x_{k+n})$ implies the label of x_k is at least $\beta + n$. Let $\zeta^{+\ell}$ be the disjoint amalgamation (in the obvious sense) of ℓ copies of ζ^- over ζ . Then if $Q_{\zeta^{+\ell}}$ is such that $\tau_\beta^-(Q_{\zeta^{+\ell}}) = \tau_\beta^-(Q_{\zeta^{+\ell}})$ and $\mathcal{M} \models Q_{\zeta^{+\ell}}(a_0, \dots, a_{k-1}, b_0, \dots, b_{\ell \cdot (n+1)-1})$ then there must be ℓ many elements among $b_0, \dots, b_{\ell \cdot (n+1)-1}$ which have label at least $\beta + n$ and which are immediate predecessors to a_{k-1} . Therefore, as ℓ was arbitrary, by β -density (b) holds.

Now suppose (b) holds. Suppose $\gamma < \beta$ and $\gamma < \text{lab}_{\mathcal{M}}(a)$. Then if $\gamma \geq \text{rank}_{\mathcal{M}}(a)$ we would have $\text{rank}_{\mathcal{M}}(a) \leq \gamma < \beta$ and so $\text{rank}_{\mathcal{M}}(a) = \text{lab}_{\mathcal{M}}(a) > \gamma$ getting a contradiction. So we must have $\gamma < \text{rank}_{\mathcal{M}}(a)$. But then β -fullness and β -true labels imply that there must be infinitely many immediate predecessors of a with label and rank γ .

Now suppose $\mathcal{M} \models Q_\zeta(a_0, \dots, a_{n-1})$ and $\zeta^- = \langle ([n+k], \leq^-), \mu^- \rangle$ is an extension of ζ . By reordering if necessary we can assume all parents precede their children in $[n+k]$. We then choose $a_n, \dots, a_{n+k-1}$ by induction where, if we have $a_0, \dots, a_{n+i-1}$ we choose a new a_{n+i} satisfying the following.

- $a_{n+i} \notin \{a_j\}_{j \in [n+i]}$.
- If $n+i <_1^- \ell$ then $a_{n+i} <_1 a_\ell$.
- If $\mu^-(n+i) < \beta$ then $\text{rank}_{\mathcal{M}}(a_{n+i}) = \mu^-(n+i)$.

- If $\mu^-(n+i) \geq \beta$ then let $(b_0, \dots, b_{z-1})$ be the longest chain among elements $\leq^- n+i$ where $\mu^-(b_j) \geq \beta$ for all $j \in [z]$. Choose a_{n+i} to be any new element with $\text{lab}_{\mathcal{M}}(a_{n+i}) \geq \beta + (z-1)$. Note that we can always find one by $(\mathfrak{A})$.

It is then immediate that there is some invariant Q_{ζ^+} where $\mathcal{M} \models Q_{\zeta^+}(a_0, \dots, a_{n+k-1})$ and $\tau_{\beta}^-(Q_{\zeta^-}) = \tau_{\beta}^-(Q_{\zeta^+})$. Therefore, as ζ and ζ^- were arbitrary, $\mathcal{M}$ satisfies β -density and (a) holds. $\square$

Putting together Theorem 3.16 and Theorem 3.18, we obtain the following.

Proposition 3.19. *Suppose $\mathcal{M}$ is a countable labeled geometric tree. Then the following are equivalent.*

- (a) $\mathcal{M}$ is α -dense for all $\alpha < \beta$ and $\tau_{\beta}^-(\mathcal{M}) = \mathcal{M}$.
- (b) *Either*
 - $\beta = 0$ and all nodes are labeled with ∞ ,
 - $\beta = \omega \cdot \delta$ is a limit ordinal and $\mathcal{M}$ is β -full, has β -true labels and has labels cut off at β , or
 - $\beta = \alpha + 1$ is a successor ordinal, $\mathcal{M}$ has true labels, $\text{range}(\text{rank}_{\mathcal{M}}) \subseteq \beta \cup \{\infty\}$, it is α -full and every element of rank ∞ has infinitely many immediate predecessors of rank ∞ .

Proof. The equivalence of (a) and (b) for $\beta = 0$ or a limit follows immediately from Theorem 3.16 and Theorem 3.18.

Now suppose $\beta = \alpha + 1$. Assume (a). We then have that $\mathcal{M}$ satisfies α -density and hence also α -fullness. We also have $\text{range}(\text{lab}_{\mathcal{M}}) \subseteq \beta \cup \{\infty\}$. If $\text{lab}_{\mathcal{M}}(a) < \alpha$ then by α -density, $\text{lab}_{\mathcal{M}}(a) = \text{rank}_{\mathcal{M}}(a)$. But then also by α -density, if $\text{lab}_{\mathcal{M}}(a) = \alpha$ it has infinitely many immediate predecessors of all ranks $< \alpha$, and so $\text{rank}_{\mathcal{M}}(a) = \alpha$. Further, if $\text{lab}_{\mathcal{M}}(a) = \infty$ then Theorem 3.18 ensures that there are infinitely many immediate predecessors b with $\text{lab}_{\mathcal{M}}(b) \geq \alpha + 1$. But as labels are cut off at $\alpha + 1$, these predecessors are labeled ∞ . Repeating this produces an infinite descending chain and so $\text{rank}_{\mathcal{M}}(a) = \infty$.

Assume (b) holds. The fact that $\mathcal{M}$ has true labels and $\text{range}(\text{rank}_{\mathcal{M}}) \subseteq \beta \cup \{\infty\}$ implies that $\tau_{\beta}^-(\mathcal{M}) = \mathcal{M}$. To show that $\mathcal{M}$ satisfies α -density we use Theorem 3.18. If for any $a \in \mathcal{M}$ we have $\alpha + n < \text{lab}_{\mathcal{M}}(a)$, then we must have $\text{rank}_{\mathcal{M}}(a) = \infty$ and its infinitely many immediate predecessors with rank $\mathcal{M}$ equal ∞ satisfy the condition in Theorem 3.18 (b). Therefore $\mathcal{M}$ satisfies α -density and (a) holds. $\square$

Now suppose $\mathbf{K}^- = \langle K_{\beta}^- \rangle_{\beta \in \omega_1}$ is a weak Knight system. If $\mathcal{M} \in K_{\beta}^-$ and $\beta > 1$ then $\mathcal{M}$ is either the one-element tree or the infinitely branching ill-founded tree in which every infinite node has infinitely many leaf children. Therefore we must have $\text{range}(\text{lab}_{\mathcal{M}}) \subseteq \{0, \infty\}$ as well. Any other structure in K^- must be in either K_0^- or K_1^- and hence also must have labels contained in $\{0, \infty\}$. In particular K^- can contain no structure with ordinal labels other than 0 and so $\mathbf{K}^-$ can't be a Knight system.

3.3. Knight System: Attempt 2. The key obstruction to τ^- occurred because of the interplay between β -density and being a fixed point of $\tau_{\beta+1}^-$. This suggests that a possible solution is to limit ourselves to limit β .

Definition 3.20. For $\beta \in \omega_1$ let $\tau_{\beta} = \tau_{\omega \cdot \beta}^-$ and let $\tau = \langle \tau_{\beta} \rangle_{\beta \in \omega_1}$.

The following is then immediate.

Lemma 3.21. $\tau = \langle \tau_{\beta} \rangle_{\beta \in \omega_1}$ is an invariant system compatible with $\mathcal{I}_{\text{tr}}$.

Theorem 3.18 suggests the following potential Knight system.

Definition 3.22. For $\beta \in \omega_1$, let K_{β} be the collection of labeled geometric trees $\mathcal{M}$ satisfying the following.

- $\mathcal{M}$ is full.
- $\mathcal{M}$ has $\omega \cdot \beta$ -true labels.
- $\mathcal{M}$ has labels cut off at $\omega \cdot \beta$.
- For all $\alpha < \beta$, $\tau_{\alpha}(\mathcal{M}) \neq \mathcal{M}$.

Now we turn our attention to maximal and eventually determined models.

Proposition 3.23. *A labeled geometric tree is maximal for $\mathbf{K}$ if and only if it is full and has true labels.*

Proof. If $\mathcal{M}$ is maximal, then $\mathcal{M} \in K_\beta$ for some β , so $\mathcal{M}$ is full by definition.

From the definition of K_β we have whenever $\mathcal{M} \in K_\beta$ that $\text{lab}_{\mathcal{M}} = \text{tr}_{\omega \cdot \beta} \circ \text{rank}_{\mathcal{M}}$. Let $\kappa(\mathcal{M}) = \min\{\delta : \text{range}(\text{rank}_{\mathcal{M}}) \subseteq \omega \cdot \delta \cup \{\infty\}\}$.

First assume that $\mathcal{M} \in K_\beta$ is full but does not have true labels. Then $\beta < \kappa(\mathcal{M})$. Let $\mathcal{N} \in K_{\kappa(\mathcal{M})}$ be the truly labeled expansion of $\mathcal{M} \upharpoonright_{\mathfrak{L}_c}$. We then have $\mathcal{N} \in K_{\kappa(\mathcal{N})}$ and $\tau_\beta(\mathcal{N}) = \tau_\beta(\mathcal{M}) = \mathcal{M}$. Therefore $\mathcal{M}$ is not maximal. Hence we have the implication from left to right.

Conversely, assume that $\mathcal{M}$ is full and has true labels. By the definition of $\kappa(\mathcal{M})$, we have $\mathcal{M} \in K_{\kappa(\mathcal{M})}$. Suppose that $\gamma > \kappa(\mathcal{M}) = \beta$, that $\mathcal{N} \in K_\gamma$, and that $\tau_\beta(\mathcal{M}) = \tau_\beta(\mathcal{N})$. Since projection preserves the underlying tree, $\mathcal{N}$ and $\mathcal{M}$ must have the same rank function. Let $a \in M$. If $\text{rank}_{\mathcal{N}}(a) \in \omega_1$ then we have $\text{rank}_{\mathcal{N}}(a) < \omega \cdot \beta$, and so $\text{lab}_{\mathcal{N}}(a) = \text{rank}_{\mathcal{N}}(a) = \text{lab}_{\mathcal{M}}(a)$. But if $\text{rank}_{\mathcal{N}}(a) = \infty$ then we must have $\text{lab}_{\mathcal{N}}(a) = \infty$ since $\text{rank}_{\mathcal{N}} \leq \text{lab}_{\mathcal{N}}$. Therefore $\mathcal{M} = \mathcal{N}$ and $\mathcal{M}$ is maximal. $\square$

Proposition 3.24. *Every full geometric tree is eventually determined with respect to $\mathbf{K}$.*

Proof. Suppose $\mathcal{M}$ is a full geometric tree and $\gamma = \min\{\beta : \text{range}(\text{rank}_{\mathcal{M}}) \subseteq \omega \cdot \beta \cup \{\infty\}\}$. Let $\mathcal{M}^* \in K_\gamma$ be the labeled geometric tree whose labelings are true and where $\mathcal{M}^*$ is a lift of $\mathcal{M}$.

Now suppose, towards a contradiction, that $\zeta > \gamma$, that $\mathcal{M}^\dagger \in K_\zeta$, and that $\mathcal{M}^\dagger$ is a lift of $\mathcal{M}$. Then $\mathcal{M}^\dagger$ must have $\omega \cdot \zeta$ -true labels and so $\mathcal{M}^\dagger = \mathcal{M}^*$. But $K_\gamma \cap K_\zeta = \emptyset$, yielding the contradiction.

Now let $\mathbf{a} \subseteq \mathcal{M}$ be any closed tuple and let $Q_{\mathbf{a}}$ be its invariant in $\mathcal{M}^*$, which captures the ranks of its elements. Then $\mathcal{M}^* \models Q_{\mathbf{a}}(\mathbf{a})$ and the second condition in Theorem 2.21 holds vacuously. Therefore $\mathbf{a}$ is eventually determined. $\square$

Putting this all together, we have the following theorem.

Theorem 3.25. *$\mathbf{K}$ is a Knight system that witnesses that Φ , the collection of full geometric trees, is a Knight class. Hence Φ has precisely $\aleph_1$ many models up to isomorphism and no perfect set of non-isomorphic models.*

Proof. The fact that $\mathbf{K}$ is a weak Knight system for Φ is immediate. It is also immediate that if $\zeta = \langle (T, \leq), \mu \rangle$ then Q_ζ is realized in any truly labeled full tree $\mathcal{M}$ with $\text{range}(\mu) \subseteq \text{range}(\text{rank}_{\mathcal{M}})$ and where the label of the root of $\mathcal{M}$ agrees with that of the root of ζ . Hence all invariants are realized in some model in $\mathbf{K}$. The fact that $\mathbf{K}$ has only countably many maximal models of any given rank follows from Theorem 3.23 and the fact that all ranks of maximal models in K_β are contained in $\omega \cdot \beta \cup \{\infty\}$. The fact that all models in Φ have maximal lifts also follows from Theorem 3.23. Finally the fact that all models in Φ are eventually determined follows from Theorem 3.24. $\square$

Hence if there were a sentence of $\mathcal{L}_{\omega_1, \omega}(\mathfrak{L}_c)$ whose countable models were precisely the full geometric trees, then that sentence would be a Knight sentence and hence a counterexample to Vaught's conjecture. However, by [Ste21], no theory of trees can be a counterexample to Vaught's conjecture. Therefore there can be no sentence of $\mathcal{L}_{\omega_1, \omega}(\mathfrak{L}_c)$ which pins down the countable full geometric trees.

4. THE COUNTEREXAMPLE

We now turn our attention to the counterexample. We will present it in three subsections. First, in Section 4.1, we will introduce the basic types of objects the counterexample is built from. This will culminate in the notion of a *template system*. Then, in Section 4.2, we will show how we can build a weak Knight system, along with a corresponding sentence of $\mathcal{L}_{\omega_1, \omega}$ from a template system. Finally, in Section 4.3, we will introduce the legal template system whose corresponding weak Knight system is an actual Knight system. This will then give us a counterexample to Vaught's conjecture.

4.1. Plans and Templates. We begin by introducing the background geometry we will want our structures to have.

Definition 4.1. A **plan** on a non-empty set A is a finitary convex geometry $\mathcal{P} \subseteq \wp_{<\omega}(A)$ in which every closed set of size at least 2 has precisely two extreme points. In other words: $\emptyset \in \mathcal{P}$, we have $\bigcup \mathcal{P} = A$, the geometry $\mathcal{P}$ is closed under intersections, every finite set is contained in some closed set, and its closure operator $\text{cl}_{\mathcal{P}}$ satisfies

- (Convexity) $(\forall B \in \mathcal{P})(\forall x, y \in A \setminus B) (x \neq y) \wedge (x \in \text{cl}_{\mathcal{P}}(B \cup \{y\})) \rightarrow (y \notin \text{cl}_{\mathcal{P}}(B \cup \{x\}))$, and
- (Extreme Points) $(\forall B \in \mathcal{P}, |B| \geq 2) |\{b \in B : B \setminus \{b\} \in \mathcal{P}\}| = 2$.

Definition 4.2. Suppose $\mathcal{P}$ is a plan on A . A **graded face** for $\mathcal{P}$ is a pair (B, j) such that $B \in \mathcal{P}$ and $1 \leq j \leq |B|$. We let $\mathfrak{GF}(\mathcal{P})$ be the collection of graded faces for $\mathcal{P}$.

If (B, j) is a graded face, we call B the **scope** and j the **grade**. We let $s_{\mathcal{P}}: \mathfrak{GF}(\mathcal{P}) \rightarrow \mathcal{P}$ be the map which takes a graded face and returns its scope and we let $g_{\mathcal{P}}: \mathfrak{GF}(\mathcal{P}) \rightarrow \omega$ be the map which takes a graded face and returns its grade. If $(B_0, j_0), (B_1, j_1) \in \mathfrak{GF}(\mathcal{P})$ then we say $(B_0, j_0) \leq (B_1, j_1)$ if $B_0 \subseteq B_1$ and $j_0 \leq j_1$.

We now introduce the combinatorial objects associated to a plan that will be used as components of invariants.

Definition 4.3. A **frame** on a plan $\mathcal{P}$ is a map $c: D \rightarrow \mathfrak{GF}(\mathcal{P})$. We call the elements of D the **cells** of the frame and define the scope of a cell d to be $s_{\mathcal{P}}(c(d))$ and the grade of a cell d to be $g_{\mathcal{P}}(c(d))$. We will abuse notation and write $s_{\mathcal{P}}(d)$ for $s_{\mathcal{P}}(c(d))$ and $g_{\mathcal{P}}(d)$ for $g_{\mathcal{P}}(c(d))$.

If $U \in \mathfrak{GF}(\mathcal{P})$ we let $D \downarrow U = \{d \in D : (s(d), g(d)) \leq U\}$. For $d \in D$ we let $D \downarrow d = D \downarrow c(d)$.

Note that the cells of a frame have no structure other than that which is given via the map into the graded faces.

Definition 4.4. Suppose c is a frame with cells D . A **label** of c is a map $\ell: D \rightarrow \{-\infty\} \cup \omega_1 \cup \{\infty\}$. A **local label** at $d \in D$ is a map $\ell_d: D \downarrow d \rightarrow \{-\infty\} \cup \omega_1 \cup \{\infty\}$.

A **row system** for c is a collection $\Sigma = (\ell_d)_{d \in D}$ such that for all $d \in D$, the label ℓ_d is a local label at d with $\text{range}(\ell_d) \subseteq \{-\infty\} \cup \omega^2$.

We are finally ready to define the type of objects which will make up our invariants.

Definition 4.5. A **template** for $\mathcal{P}$ is a triple $t = (c, \Sigma, p)$ where c is a frame for $\mathcal{P}$, where $\Sigma = (\ell_d)_{d \in D}$ is a row system for c , and where p is a label for c .

We define the **range** of the template t to be $\text{range}(t) = \text{range}(p)$. For $\alpha \in \omega_1$ we say that t is an **α -template** if $\text{range}(t) \subseteq \{-\infty\} \cup \alpha \cup \{\infty\}$.

There are two types of equivalence of templates which will be important: isomorphism and renaming.

Definition 4.6. Suppose $\iota: D_0 \rightarrow D_1$ is a bijection.

- Suppose $c: D_1 \rightarrow \mathfrak{GF}(\mathcal{P})$ is a frame. We say that $c \circ \iota$ is the **renaming of c by ι** .
- Suppose $\ell: D_1 \rightarrow \{-\infty\} \cup \omega_1 \cup \{\infty\}$ is a label. We say that $\ell \circ \iota$ is a **renaming of ℓ by ι** .
- Suppose $\Sigma = (\ell_d)_{d \in D_1}$ is a row system. We define $(\ell_{\iota(d)} \circ \iota \upharpoonright_{D_0 \downarrow d})_{d \in D_0}$ to be a **renaming of Σ by ι** .
- Suppose $t = (c, \Sigma, p)$ is a template and c', Σ' are the renamings by ι of c and Σ respectively. We call $(c', \Sigma', p \circ \iota)$ the **renaming of t by ι** .

For templates t_0, t_1 , we write $t_0 \simeq t_1$ when t_1 is a renaming of t_0 by some map.

As the cells are intended to have no internal structure, the specific sets representing cells are not relevant for the notion of a template, and we think of renamings of templates as being “the same”.

Definition 4.7. Suppose $i: A \rightarrow A^*$ is a bijection and $\mathcal{P}^* = \{i[B] : B \in \mathcal{P}\}$.

- We say that i is an **isomorphism from $(A, \mathcal{P})$ to $(A^*, \mathcal{P}^*)$** .
- Suppose $c: D \rightarrow \mathfrak{GF}(\mathcal{P})$ and $c^*: D \rightarrow \mathfrak{GF}(\mathcal{P}^*)$ are frames. We say that i is an **isomorphism from c to c^*** if $c^* = \hat{i} \circ c$ where $\hat{i}: \mathfrak{GF}(\mathcal{P}) \rightarrow \mathfrak{GF}(\mathcal{P}^*)$ is the map that takes (B, j) to $(i[B], j)$.

- Suppose $t = (c, \Sigma, p)$ and $t^* = (c^*, \Sigma, p)$ are templates. We say that i is an **isomorphism from t to t^*** if it is an isomorphism from c to c^* .

Note that isomorphisms only deal with reindexing of the underlying set of the plan, not with renaming the cells.

We now define two different ways we can restrict a template.

Definition 4.8. Suppose $\mathcal{P}$ is a plan on A with $\emptyset \neq A_0 \subseteq A$ and suppose A_0 is closed, i.e., $A_0 \in \mathcal{P}$.

- We define $\mathcal{P} \upharpoonright_{A_0} = \{X \in \mathcal{P} : X \subseteq A_0\}$. We call $\mathcal{P} \upharpoonright_{A_0}$ the **restriction** of $\mathcal{P}$ to A_0 .
- Suppose $c : D \rightarrow \mathfrak{GF}(\mathcal{P})$ is a frame. We define $D_0 = D \upharpoonright_{A_0} = \{d \in D : s_{\mathcal{P}}(d) \subseteq A_0\}$ and $c \upharpoonright_{A_0} : D_0 \rightarrow \mathfrak{GF}(\mathcal{P} \upharpoonright_{A_0})$ to be the map where $c \upharpoonright_{A_0} = c \upharpoonright_{D_0}$.
- Suppose $\ell : D \rightarrow \{-\infty\} \cup \omega_1 \cup \{\infty\}$ is a label. We let $\ell \upharpoonright_{A_0} = \ell \upharpoonright_{D_0}$. For $d \in D_0$, if ℓ_d is a local label we let $\ell_d \upharpoonright_{A_0} = \ell_d \upharpoonright_{D_0 \downarrow d}$.
- Suppose $\Sigma = (\ell_d)_{d \in D}$ is a row system. We define $\Sigma \upharpoonright_{A_0} = (\ell_d \upharpoonright_{A_0})_{d \in D_0}$.
- Suppose $t = (c, \Sigma, p)$ is a template. We define $t \upharpoonright_{A_0} = (c \upharpoonright_{A_0}, \Sigma \upharpoonright_{A_0}, p \upharpoonright_{A_0})$.
- Suppose t is a template on A_t . We say that a template u **extends** t if A_t is closed in the underlying plan of u and $u \upharpoonright_{A_t} = t$.

Note that the ordering between graded faces does not change when the plan is restricted. Specifically, if $A_0 \subseteq A$ is closed and $(B_0, j_0), (B_1, j_1) \in \mathfrak{GF}(\mathcal{P}) \cap \mathfrak{GF}(\mathcal{P} \upharpoonright_{A_0})$, then whether or not (B_0, j_0) is less than (B_1, j_1) does not depend on whether they are considered as elements of $\mathfrak{GF}(\mathcal{P})$ or $\mathfrak{GF}(\mathcal{P} \upharpoonright_{A_0})$.

We now turn to the notion of a β -truncation.

Definition 4.9. We extend the β -truncation function tr_{β} from Theorem 3.12 so that $\text{tr}_{\beta}(-\infty) = -\infty$.

Definition 4.10. Suppose $\beta \in \omega_1$. We define the β -**truncation** of a template $t = (c, \Sigma, p)$ to be $(c, \Sigma, \text{tr}_{\beta} \circ p)$. We denote it by $\text{tr}_{\beta}(t)$.

Just as in the case of full trees, the β -truncation of a template takes all values of the label which are $\geq \beta$ and replaces them with ∞ .

Definition 4.11. Suppose T is a collection of templates which satisfy the following whenever $t = (c, \Sigma, p) \in T$ with underlying finite set A .

- Both the underlying set of t and the cells of t are finite subsets of ω .
- If $i : A \rightarrow B$ is a bijection with $B \subseteq \omega$ then $(i \circ c, \Sigma, p) \in T$.
- If $\emptyset \neq A_0 \subseteq A$ is closed then $t \upharpoonright_{A_0} \in T$.
- For all $\beta \in \omega_1$, we have $\text{tr}_{\omega + \omega \cdot \beta}(t) \in T$.
- For all $\beta \in \omega_1$, there is a $t_{\beta} \in T$ such that $\text{tr}_{\omega + \omega \cdot \beta}(t_{\beta}) \neq t_{\beta}$.
- For all $n \in \omega$, there is a $t_n \in T$ whose underlying set has size at least n .
- For every $\beta \in \omega_1$, letting $\lambda = \omega + \omega \cdot \beta$, every $\delta < \lambda$, and all templates $t, u \in T$, if there is a template u' with $u' \simeq u$ such that u' extends $\text{tr}_{\lambda}(t)$, then there are a template $v \in T$ on the same underlying set as u and a template v' with $v' \simeq v$ such that v' extends t and $\text{tr}_{\delta}(v) \simeq \text{tr}_{\delta}(u)$.

We say that T^* is a **template system** if the following hold.

- $T^* \subseteq T$.
- If $t \in T^*$ then the underlying set of t and the cells of t are initial segments of ω .
- If $t \in T$ has underlying set $[n]$ then there is some $t^* \in T^*$ which is a renaming of t .
- If $t, t' \in T^*$ and $t \simeq t'$ then $t = t'$.

In Theorem 4.11, condition (a) tells us that the underlying set as well as the cells of anything in T are finite subsets of ω . In this way, we are able to bound the size of the collection T . Condition (b) tells us that T is closed under isomorphisms of templates and conditions (c) and (d) tell us that T is closed under the restriction to smaller sets as well as appropriate truncation. Conditions (e) and (f) tell us that there are *enough* templates in T in the sense that there are arbitrarily *complex* templates and templates whose underlying set is arbitrarily large. Condition (g) tells us that there are enough templates to allow the density axioms to function the way we would like.

There is one challenge with regarding classes that satisfy conditions (a) – (g) as the invariants, and that is the fact that information can be encoded in both the structure of the specific underlying sets and cells chosen. To get around this, we only consider templates whose underlying set and cells are each initial segments of ω , and we also only allow one template per $\simeq$ -class.

4.2. Knight Systems. We now show how templates can be used as invariants and how to use a template system to build a weak Knight system.

Definition 4.12. Suppose $\mathfrak{L}_a = \{P_n\}_{n \geq 1}$ where each P_n has arity n . We say that an $\mathfrak{L}_a$ -structure $\mathcal{M}$ is a **plan** if $\mathcal{M} \models \text{Th}_{\text{Geo}}(\mathfrak{L}_a)$ and $\{\emptyset\} \cup \bigcup_{n \geq 1} \{\{a_i\}_{i \in [n]} : \mathcal{M} \models P_n(a_0, \dots, a_{n-1})\}$ is a plan.

We now define the theory describing those geometries which are plans.

Definition 4.13. Let $\text{Th}_{\text{pl}} \subseteq \mathcal{L}_{\omega_1, \omega}(\mathfrak{L}_a)$ consist of the following sentences.

- $\text{Th}_{\text{Geo}}(\mathfrak{L}_a)$.
- (Singletons) $(\forall x) P_1(x)$.
- (Convexity)

$$(\forall x_0, \dots, x_{n-1})(\forall y, z) \left[P_n(x_0, \dots, x_{n-1}) \wedge (y \neq z) \wedge \bigwedge_{i \in [n]} [(x_i \neq y) \wedge (x_i \neq z)] \right] \\ \rightarrow [y \in \text{cl}(\{x_i\}_{i \in [n]} \cup \{z\}) \rightarrow z \notin \text{cl}(\{x_i\}_{i \in [n]} \cup \{y\})],$$

where $a \in \text{cl}(\{b_i\}_{i \in [n]})$ is the infinitary formula which says that there is a finite set $\{e_i\}_{i \in [m]}$ such that $\{b_i\}_{i \in [n]} \cup \{a\} \subseteq \{e_i\}_{i \in [m]}$, that $P_m(e_0, \dots, e_{m-1})$ holds, and that for any $\{i_j\}_{j \in [k]}$ if $P_k(e_{i_0}, \dots, e_{i_{k-1}})$ and b_r is among $\{e_{i_j}\}_{j \in [k]}$ for each $r \in [n]$ then a is among $\{e_{i_j}\}_{j \in [k]}$.

- (Extreme Points) For all $n \geq 3$, the following two sentences hold:

$$(\forall x_0, \dots, x_{n-1}) P_n(x_0, \dots, x_{n-1}) \\ \rightarrow \bigvee_{i < j \in [n]} [P_{n-1}(x_0, \dots, x_{i-1}, x_{i+1}, \dots, x_{n-1}) \wedge P_{n-1}(x_0, \dots, x_{j-1}, x_{j+1}, \dots, x_{n-1})]; \\ (\forall x_0, \dots, x_{n-1}) P_n(x_0, \dots, x_{n-1}) \\ \rightarrow \neg [P_{n-1}(x_1, \dots, x_{n-1}) \wedge P_{n-1}(x_0, x_2, \dots, x_{n-1}) \wedge P_{n-1}(x_0, x_1, x_3, \dots, x_{n-1})].$$

In Theorem 4.13, the (Singletons) sentence ensures that all singletons are closed, the (Convexity) sentences ensure that the geometry is convex and the (Extreme Points) sentences ensure that every closed set of size at least 2 has precisely 2 extreme points.

The following is then immediate.

Lemma 4.14. An $\mathfrak{L}_a$ -structure $\mathcal{M}$ is a plan if and only if $\mathcal{M} \models \text{Th}_{\text{pl}}$.

We now introduce the languages $\mathfrak{L}_b$ and the corresponding invariant diagram and invariant system.

Definition 4.15. Suppose T is a template system. We define $\mathfrak{L}_b[T] = \{Q_t\}_{t \in T}$ where Q_t has arity n when the underlying set of t is $[n]$.

We now can introduce the invariant system. Note that we are choosing to have the projections τ_β project down to templates with labels in $\{-\infty\} \cup \omega + \omega \cdot \beta \cup \{\infty\}$. See the discussion at the end of Section 3.2 for motivation for why this is necessary.

Definition 4.16. Suppose $\beta \in \omega_1$. For each $t \in T$, we let $\tau_\beta(Q_t) = Q_{t^*}$, where t^* is the unique element of T satisfying $t^* \simeq \text{tr}_{\omega + \omega \cdot \beta}(t)$. We let $\boldsymbol{\tau} = \langle \tau_\beta \rangle_{\beta \in \omega_1}$.

Note that as $\text{tr}_{\omega + \omega \cdot \alpha} \circ \text{tr}_{\omega + \omega \cdot \beta} = \text{tr}_{\omega + \omega \cdot \beta} \circ \text{tr}_{\omega + \omega \cdot \alpha} = \text{tr}_{\omega + \omega \cdot \alpha}$ whenever $\alpha \leq \beta$, we have that $\boldsymbol{\tau}$ is an invariant system.

We now introduce our (countable) language $\mathfrak{L}_c[T]$. For technical reasons having to do with the abstract machinery of Knight systems we want $\mathfrak{L}_b[T]$ and $\mathfrak{L}_c[T]$ to be disjoint. However, the intuition is that the base language should be those invariants which have rank 0.

Definition 4.17. We let $\mathfrak{L}_c[T] = \mathfrak{L}_a \cup \{R_t : t \in T \text{ and } \tau_0(Q_t) = Q_t\}$, where each R_t is a new relation symbol of the same arity as Q_t .

We are now ready to define our invariant diagram.

Definition 4.18. Let $\mathcal{I}_{\text{temp}}[T]$ be the invariant diagram such that whenever $t = (c, \Sigma, p)$ is a template with underlying set $[n]$ then $(Q_t(x_0, \dots, x_{n-1}), R(x_{i_0}, \dots, x_{i_{k-1}})) \in \mathcal{I}_{\text{temp}}[T]$ precisely when one of the following holds.

- R is P_k , $i_0, \dots, i_{k-1}$ are distinct elements and $\{i_j\}_{j \in [k]}$ is closed in the geometry of t .
- R is Q_{t^*} where the underlying set of t^* is $[k]$ and the map $i_j \mapsto j$ is an isomorphism between $t^\dagger$ and t^* where $t^\dagger \simeq t \upharpoonright_{\{i_j\}_{j \in [k]}}$.
- R is R_{t^-} where $(Q_t(x_0, \dots, x_{n-1}), Q_{t^+}(x_{i_0}, \dots, x_{i_{k-1}})) \in \mathcal{I}_{\text{temp}}[T]$ and $t^- \simeq \text{tr}_\omega(t^+)$.

We call a structure compatible with $\mathcal{I}_{\text{temp}}[T]$ a **T -structure**.

The following is immediate.

Lemma 4.19. $\mathcal{I}_{\text{temp}}[T]$ and τ are compatible.

Before we introduce our weak Knight system, we will need a more fine-grained notion of density than what will be guaranteed by the invariant system. This is analogous to how, in the case of trees, we wanted trees in K_β to have density at all levels below $\omega \cdot \beta$, not just the limit levels that would be forced by the conditions on being in a weak Knight system.

Definition 4.20. We let $D[\beta]$ be the collection of axioms

$$(\forall \mathbf{x}) \left(Q_t(\mathbf{x}) \rightarrow (\exists \mathbf{y}) \bigvee_{\substack{s \in T \\ \text{tr}_\beta(s) \simeq \text{tr}_\beta(u)}} Q_s(\mathbf{xy}) \right)$$

for each $(Q_u(\mathbf{xy}), Q_t(\mathbf{x})) \in \mathcal{I}_{\text{temp}}[T]$ with $u, t \in T$.

We now define the weak Knight system associated to T .

Definition 4.21. For $\beta \in \omega_1$, let $K_\beta[T]$ be the collection of $(\mathfrak{L}_b[T] \cup \mathfrak{L}_c[T])$ -structures $\mathcal{M}$ such that

- $\mathcal{M}$ is compatible with $\mathcal{I}_{\text{temp}}[T]$,
- $\tau_\beta(\mathcal{M}) = \mathcal{M}$, and
- $\mathcal{M} \models D[\alpha]$ for all $\alpha < \omega + \omega \cdot \beta$.

We let $\mathbf{K}[T] = \langle K_\beta[T] \rangle_{\beta \in \omega_1}$, let $K[T] = \bigcup_{\beta \in \omega_1} K_\beta[T]$ and let $K_{-\infty}[T] = \{\mathcal{M} \upharpoonright_{\mathfrak{L}_c[T]} : \mathcal{M} \in K[T]\}$.

We now give the countable theory of $\mathcal{L}_{\omega_1, \omega}(\mathfrak{L}_c)$ which captures the base of $\mathbf{K}[T]$.

Definition 4.22. Let $\sigma[T] \in \mathcal{L}_{\omega_1, \omega}(\mathfrak{L}_c)$ be the (countable) conjunction of the following sentences where $\mathcal{R}[T] = \mathfrak{L}_c[T] \setminus \mathfrak{L}_a$.

- $\bigwedge \text{Th}_{\text{pl}}$.
- For all $n \geq 1$,
 - for all $P_n \in \mathfrak{L}_a^n$,

$$(\forall x_0, \dots, x_{n-1}) P_n(x_0, \dots, x_{n-1}) \leftrightarrow \bigvee_{R_n \in \mathcal{R}[T]^n} R_n(x_0, \dots, x_{n-1});$$

- for all distinct $R_n, R'_n \in \mathcal{R}[T]^n$,

$$(\forall x_0, \dots, x_{n-1}) R_n(x_0, \dots, x_{n-1}) \rightarrow \neg R'_n(x_0, \dots, x_{n-1});$$

- for any $k \geq 1$, any $R \in \mathcal{R}[T]^n$, any $S \in \mathfrak{L}_c[T]^k$, and any sequence $\langle i_j \rangle_{j \in [k]} \subseteq [n]$,

$$(\forall x_0, \dots, x_{n-1}) (\forall y_0, \dots, y_{n-1}) [R(x_0, \dots, x_{n-1}) \wedge R(y_0, \dots, y_{n-1})] \\ \rightarrow [S(x_{i_0}, \dots, x_{i_{k-1}}) \leftrightarrow S(y_{i_0}, \dots, y_{i_{k-1}})].$$

- (c) If $(Q_t(\mathbf{x}), \psi(\mathbf{y})) \in \mathcal{I}_{\text{temp}}[T]$, where $R_t \in \mathcal{R}[T]$ and $\psi(\mathbf{y})$ is an atomic or negated atomic $\mathfrak{L}_c[T]$ -formula,

$$(\forall \mathbf{x}) [R_t(\mathbf{x}) \rightarrow \psi(\mathbf{y})].$$

- (d) For all $n \in \omega$, for all $t, t^* \in T$, if $R_t \in \mathfrak{L}_c[T]$ and $R_{t^*} \in \mathcal{R}[T]$ with $(Q_t(\mathbf{xy}), Q_{t^*}(\mathbf{x})) \in \mathcal{I}_{\text{temp}}[T]$, then

$$(\forall \mathbf{x}) \left[R_{t^*}(\mathbf{x}) \rightarrow (\exists \mathbf{y}) \bigvee_{\substack{s \in T \\ R_s \in \mathfrak{L}_c[T] \\ \text{tr}_n(s) \simeq \text{tr}_n(t)}} R_s(\mathbf{xy}) \right].$$

The intuition is that the relation symbols in $\mathfrak{L}_c[T] \setminus \mathfrak{L}_a$ are meant to capture the relation symbols in $\mathfrak{L}_b[T]$ of rank 0. As such, Theorem 4.22 (a) describes the geometry, Theorem 4.22 (b) and (c) provide the remaining conditions for structured geometry and enforce compatibility with $\mathcal{I}_{\text{temp}}[T]$, and Theorem 4.22 (d) is equivalent to density.

We then have the following.

Theorem 4.23. *$K[T]$ is a weak Knight system whose class of $\mathfrak{L}_c$ -reducts is the collection of models of $\sigma[T]$.*

Proof. All the conditions of $K[T]$ being a weak Knight system are immediately satisfied, except for the following claim.

Claim 4.24. *For any $\alpha < \beta$ and $\mathcal{M} \in K_\beta$, $\tau_\alpha(\mathcal{M}) \in K_\alpha$.*

Proof. Write $\lambda_\alpha = \omega + \omega \cdot \alpha$. Let $\mathcal{N} = \tau_\alpha(\mathcal{M})$. Idempotence of τ_α gives $\tau_\alpha(\mathcal{N}) = \mathcal{N}$. Further, as $\mathcal{M}$ is compatible with $\mathcal{I}_{\text{temp}}[T]$ and compatibility with $\mathcal{I}_{\text{temp}}[T]$ is preserved by projection, we have that $\mathcal{N}$ is compatible with $\mathcal{I}_{\text{temp}}[T]$.

It remains to show $\mathcal{N} \models D[\delta]$ for all $\delta < \lambda_\alpha$. Fix such a δ and suppose $\mathcal{N} \models Q_s(\mathbf{a})$ and u is a template extending s . We must find an extension of $\mathbf{a}$ in $\mathcal{N}$ whose template agrees with u after truncating to δ .

Let t be the template such that $\mathcal{M} \models Q_t(\mathbf{a})$. We then have $s \simeq \text{tr}_{\lambda_\alpha}(t)$. By Theorem 4.11 (g), there is a template v on the same underlying set as u whose restriction to the underlying set of t is a renaming of t and such that $\text{tr}_\delta(v) \simeq \text{tr}_\delta(u)$.

Since $\delta < \lambda_\alpha < \omega + \omega \cdot \beta$, we have $\mathcal{M} \models D[\delta]$. Applying this to the extension v of t we get a tuple $\mathbf{b} \subseteq \mathcal{M}$ and template w such that $\mathcal{M} \models Q_w(\mathbf{ab})$ and $\text{tr}_\delta(w) \simeq \text{tr}_\delta(v)$. But in $\mathcal{N}$, the tuple $\mathbf{ab}$ realizes the invariant associated to $\text{tr}_{\lambda_\alpha}(w)$. So $\text{tr}_\delta(\text{tr}_{\lambda_\alpha}(w)) = \text{tr}_\delta(w) \simeq \text{tr}_\delta(u)$. This is the required witness to $D[\delta]$. As δ was arbitrary less than $\omega + \omega \cdot \alpha$, this implies that $\tau_\alpha(\mathcal{M}) \in K_\alpha$. $\square$

Therefore K is a weak Knight system and it suffices to show that its base is the collection of models of $\sigma[T]$, i.e., $\{\mathcal{M} : \mathcal{M} \models \sigma[T]\} = \{\mathcal{M} \upharpoonright_{\mathfrak{L}_c[T]} : \mathcal{M} \in K_0\}$. But this follows from the fact that for every invariant Q_t fixed by τ_0 there is an equivalent relation symbol $R_t \in \mathfrak{L}_c[T]$ and up to that equivalence the models of $\sigma[T]$ are precisely the structures in K_0 . $\square$

Putting this all together we have the following.

Theorem 4.25. *If $K[T]$ is a Knight system then $\sigma[T]$ is a Knight sentence and hence a counterexample to Vaught's conjecture for $\mathcal{L}_{\omega_1, \omega}$.*

Proof. This is immediate from Theorems 2.27 and 4.23. $\square$

4.3. Legal Templates. We have now reduced the problem of finding a counterexample to Vaught's conjecture for $\mathcal{L}_{\omega_1, \omega}$ to the problem of finding a template system such that $K[T]$ is a Knight system. Finding such a template system required a great deal of work using coding agents to interact with LLMs, Lean, and SMT and Horn clause solvers. The first example we found was substantially more complicated; here we present a considerably simplified such example.

We will need several definitions, the first of which are the “visibility maps” from $\{-\infty\} \cup \omega_1 \cup \{\infty\}$ to $\{-\infty\} \cup \omega_1 \cup \{\infty\}$.

Definition 4.26. Let $i, k \in \omega$. We let $\text{vis}_{k:i}: \{-\infty\} \cup \omega_1 \cup \{\infty\} \rightarrow \{-\infty\} \cup \omega_1 \cup \{\infty\}$ be the map where

- $\text{vis}_{k:i}(-\infty) = -\infty$ and $\text{vis}_{k:i}(\infty) = \infty$, and
- for $\alpha \in \omega_1$ and $n \in \omega$,

$$\text{vis}_{k:i}(\omega \cdot \alpha + n) = \begin{cases} \omega \cdot \alpha + i & \text{if } n < k, \\ \omega \cdot \alpha + n & \text{otherwise.} \end{cases}$$

We say that $x \in \{-\infty\} \cup \omega_1 \cup \{\infty\}$ is **self-visible at k** if $\text{vis}_{k:k}(x) = x$.

Definition 4.27. Suppose r is a frame with cells D and u, v are labelings of r . We write $u \Rightarrow v$ if there are functions $f: \omega \rightarrow \{-\infty\} \cup \omega_1 \cup \{\infty\}$ and $h: \{-\infty\} \cup \omega_1 \cup \{\infty\} \rightarrow \{-\infty\} \cup \omega_1 \cup \{\infty\}$ which satisfy the following.

- f is antitone, i.e., for all $a \leq b \in \omega$, $f(a) \geq f(b)$.
- For all $k \in \omega$, the value $f(k)$ is self-visible at k .
- h is non-decreasing.
- $h(-\infty) = -\infty$.
- For all $i \leq k \in \omega$ and x with $h(x) \leq f(k)$, we have

$$h(\text{vis}_{k:i}(x)) = \text{vis}_{k:i}(h(x)).$$

- For all $d \in D$, we have

$$v(d) = \min\{h(u(d)), f(g(d))\},$$

where $g(d)$ is the grade of d .

We now give the defining property of our template system.

Definition 4.28. Suppose $t = (c, \Sigma, p)$ is a template with $\Sigma = (\ell_d)_{d \in D}$ and g, s are the grade and scope functions. For U a graded face a **local labeling** $r: D \downarrow U \rightarrow \{-\infty\} \cup \omega_1 \cup \{\infty\}$ is **lawful** with respect to Σ if it satisfies the following.

- (Orderliness) For all $d \in D \downarrow U$, the value $r(d)$ is self-visible at $g(d)$.
- (Locality) For all $d \in D \downarrow U$, we have $\ell_d \Rightarrow r_d$, where $r_d: D \downarrow d \rightarrow \{-\infty\} \cup \omega_1 \cup \{\infty\}$ is defined by $r_d(e) = \min\{r(e), r(d)\}$.
- (Availability) For all $d, e \in D \downarrow U$, if $s(d) \subseteq s(e)$ and $g(d) = g(e)$ then there is an $e' \in D$ with $c(e) = c(e')$ and $r(d) \leq r(e')$.

We define a **lawful labeling** with respect to Σ in the obvious analogous way. We say a (local) labeling is lawful with respect to t when it is lawful with respect to Σ .

For $\gamma \in \{-\infty\} \cup \omega_1 \cup \{\infty\}$ and q a lawful local labeling (with respect to t) let

$$B_\gamma(q) = \{r : r \text{ is lawful and } \min\{r, \gamma\} = \min\{q, \gamma\}\}.$$

We say the rows of Σ are **bountiful** if for every graded face U , lawful local map q with $\text{dom}(q) = D \downarrow U$, $\gamma \in \{-\infty\} \cup \omega_1 \cup \{\infty\}$ which is self-visible at $g(U)$ and graded face $V \leq U$ we have

$$\{r \upharpoonright_{D \downarrow V} : r \in B_\gamma(q)\} = B_\gamma(q \upharpoonright_{D \downarrow V}).$$

Definition 4.29. Suppose $t = (c, \Sigma, p)$ is a template over the plan $(A, \mathcal{P})$ with $\Sigma = (\ell_d)_{d \in D}$. We say t is **legal** if it satisfies the following.

- $\text{range}(c) = \mathfrak{GF}(\mathcal{P})$, i.e., every graded face has a cell.
- p is lawful with respect to Σ .
- For each $d \in D$, ℓ_d is lawful.
- The rows of $(\ell_d)_{d \in D}$ are bountiful.
- For each d , $\text{range}(\ell_d) \subseteq \{-\infty\} \cup \{\omega \cdot n + i : n \in \omega, i \leq g(d) + 1\}$.

We then have the following proposition.

Proposition 4.30. *There is a template system L such that*

- *all templates in L are finite legal templates, and*

- all finite legal templates are isomorphic to renamings of templates in L .

Proof. Let T be the collection of legal templates whose underlying sets and cells are finite subsets of ω . It suffices to verify Theorem 4.11 (a) – (g) as then we can let L choose one representative of each cell-renaming class on each underlying set $[n]$ with cell set an initial segment of ω .

First, Theorem 4.11 (a) – (c), (e), (f) are straightforward. To see (d) suppose $t = (c, \Sigma, p)$ is legal and let $\lambda = \omega + \omega \cdot \beta$. We must show $(c, \Sigma, \text{tr}_\lambda \circ p)$ is legal. Note tr_λ is non-decreasing, preserves minima and commutes with visibility, i.e. $\text{tr}_\lambda \circ \text{vis}_{k:i} = \text{vis}_{k:i} \circ \text{tr}_\lambda$.

Now suppose $u \Rightarrow v$ and f, h witness this fact. Let $f' = \text{tr}_\lambda \circ f$ and $h' = \text{tr}_\lambda \circ h$. It is then straightforward to check that f', h' witness that $u \Rightarrow \text{tr}_\lambda \circ v$.

Finally we need to verify that $\text{tr}_\lambda \circ p$ is lawful. But each of orderliness, locality, and availability follows straightforwardly from the above. Therefore $\text{tr}_\lambda(t)$ is legal, proving Theorem 4.11 (d).

To see Theorem 4.11 (g) suppose t, u are legal templates and some renaming of u extends $\text{tr}_\lambda(t)$ where $\lambda = \omega + \omega \cdot \beta$. Since legality is preserved by cell renaming, we may assume without loss of generality that u literally extends $\text{tr}_\lambda(t)$. Write $A_t \subseteq A_u$ for their underlying sets and let p_t, p_u be the corresponding labelings. We also identify the cells of t with the corresponding cells of u . Let $V_t = (A_t, |A_t|)$, $V_u = (A_u, |A_u|)$.

Choose $\delta < \lambda$. If $A_u = A_t$ take $v = t$. Otherwise $A_t \subsetneq A_u$. Choose γ where $\delta \leq \gamma < \lambda$ and γ is self-visible at $|A_u|$. Such a γ exists as λ is a limit ordinal. Since u extends $\text{tr}_\lambda(t)$, we have

$$\min\{p_t, \gamma\} = \min\{p_u \upharpoonright_{D_u \downarrow V_t}, \gamma\},$$

where D_u is the set of cells of u .

Both labelings are lawful with respect to rows restricted to $D_u \downarrow V_t$. Thus $p_t \in B_\gamma(p_u \upharpoonright_{D_u \downarrow V_t})$. So, by bountifulness, there is an $r \in B_\gamma(p_u)$ such that $r \upharpoonright_{D_u \downarrow V_t} = p_t$.

Let $v = (c_u, \Sigma_u, r)$ (where $u = (c_u, \Sigma_u, p_u)$). Since r is lawful and the frame and rows are those of u , v is legal. It extends t and has the same underlying set as u .

Finally, because $r \in B_\gamma(p_u)$ we have $\min\{r, \gamma\} = \min\{p_u, \gamma\}$. So, because $\delta \leq \gamma$, $\text{tr}_\delta \circ r = \text{tr}_\delta \circ p_u$. Therefore $\text{tr}_\delta(v) = \text{tr}_\delta(u)$ as required. So Theorem 4.11 (g) holds. $\square$

Note the collection of legal templates does not itself form a template system, as template systems by design have only a single element out of each cell-renaming class. Between Theorem 4.23 and Theorem 4.30, the class of models of $\sigma[L]$ is the base of a weak Knight system $\mathbf{K}[L]$.

All that is left to show $\sigma[L]$ is a counterexample to Vaught's conjecture is to show that $\mathbf{K}[L]$ is in fact a Knight system. This proof, which is substantial, has been formalized. We will now give a rough outline of the approach. Specifically we will mention three propositions without proof and show how to prove from them that $\mathbf{K}[L]$ is a Knight system.

Definition 4.31. Suppose $t \in L$ has underlying set $[n]$. We define a **chart** on t to be a finite structure $\mathcal{M}$ of size n such that for some enumeration $\{a_i\}_{i \in [n]}$ of $\mathcal{M}$ we have the following for all $R \in \mathfrak{L}_b[L]^k \cup \mathfrak{L}_c[L]^k$ and all $\langle i_j \rangle_{j \in [k]} \subseteq [n]$:

$$\mathcal{M} \models R(a_{i_0}, \dots, a_{i_{k-1}}) \text{ if and only if } (Q_t(x_0, \dots, x_{n-1}), R(x_{i_0}, \dots, x_{i_{k-1}})) \in \mathcal{I}_{\text{temp}}[L].$$

For $\beta \in \omega_1$, let $C_\beta[L]$ be the category of charts on t such that $\text{tr}_{\omega+\omega \cdot \beta}(t) = t$, with morphisms the embeddings between them.

A chart is then a finite structure determined by the invariant of an enumeration of its whole domain.

Proposition 4.32. For all $\beta \in \omega_1$ the category $C_\beta[L]$ satisfies the following.

- (Hereditary Property) If $\mathcal{M} \in C_\beta[L]$ and $\emptyset \neq M_0 \subseteq M$ with M_0 closed then $\mathcal{M} \upharpoonright_{M_0} \in C_\beta[L]$.
- (Joint Embedding Property) If $\mathcal{M}, \mathcal{N} \in C_\beta[L]$ then there is a $\mathcal{R} \in C_\beta[L]$ with embeddings $i: \mathcal{M} \rightarrow \mathcal{R}$ and $j: \mathcal{N} \rightarrow \mathcal{R}$.
- (Amalgamation Property) If $\mathcal{M}, \mathcal{N}, \mathcal{R} \in C_\beta[L]$ with $\mathcal{M} \subseteq \mathcal{N}$ and $\mathcal{M} \subseteq \mathcal{R}$ then there is a $\mathcal{Q} \in C_\beta[L]$ and embeddings $i: \mathcal{N} \rightarrow \mathcal{Q}$ and $j: \mathcal{R} \rightarrow \mathcal{Q}$ such that $j \upharpoonright_{\mathcal{M}} = i \upharpoonright_{\mathcal{M}}$.

Lemma 4.33. For each $\beta \in \omega_1$ there is then a structure $\mathcal{F}_\beta \in K_\beta$ satisfying the following.

- Every chart in $C_\beta[L]$ embeds into $\mathcal{F}_\beta$.
- Whenever
 - $Q_s, Q_t \in \mathfrak{L}_b[L][\tau, \beta]$,
 - $(Q_t(\mathbf{x}\mathbf{y}), Q_s(\mathbf{x})) \in \mathcal{I}_{\text{temp}}[L]$,
 - $\mathcal{F}_\beta \models Q_s(\mathbf{a})$,
 there is a $\mathbf{b} \subseteq \mathcal{F}_\beta$ such that $\mathcal{F}_\beta \models Q_t(\mathbf{ab})$.

We call $\mathcal{F}_\beta$ the β -**Fraïssé limit** of $\mathbf{K}$.

Proof. Since $C_\beta[L]$ has only countably many isomorphism types the existence of $\mathcal{F}_\beta$ follows immediately from Theorem 4.32 and the standard Fraïssé construction. The fact that $\mathcal{F}_\beta \in K_\beta$ follows from the fact that the extension property of $\mathcal{F}_\beta$ implies $\mathcal{F}_\beta$ satisfies $D[\delta]$ for all $\delta < \omega + \omega \cdot \beta$. $\square$

We now describe three key properties of the lifts of a specific $\mathfrak{L}_c[L]$ -structure.

Proposition 4.34. *For every $\beta \in \omega_1$ and every $\mathcal{M} \models \sigma[L]$ the following hold.*

- (Boundedness) The set $\{\text{rank}_\tau(\mathcal{N}) : \mathcal{N} \in K \text{ and } \mathcal{N} \upharpoonright_{\mathfrak{L}_c[L]} = \mathcal{M}\}$ is countable.
- (Reconstruction) If $\mathcal{N}, \mathcal{R} \in K_\beta$ and $\mathcal{N} \upharpoonright_{\mathfrak{L}_c[L]} = \mathcal{R} \upharpoonright_{\mathfrak{L}_c[L]} = \mathcal{M}$ then $\mathcal{N} = \mathcal{R}$.
- (Limits) If β is a limit ordinal and for all $\alpha < \beta$ there is a $\mathcal{N}_\alpha \in K_\alpha$ with $\mathcal{N}_\alpha \upharpoonright_{\mathfrak{L}_c[L]} = \mathcal{M}$ then there is a $\mathcal{N}_\beta \in K_\beta$ such that $\mathcal{N}_\beta \upharpoonright_{\mathfrak{L}_c[L]} = \mathcal{M}$.

Boundedness tells us that no structure in the base has lifts of arbitrarily large ranks. Reconstruction tells us that lifts of a structure at a given level must be unique. Limits tell us that if β is a limit and a structure has lifts of all levels $\alpha < \beta$ then it must have a lift of level β . It is worth observing that by Theorem 3.23 all three of these conditions also hold of the collection of full trees and their corresponding Knight system.

Proposition 4.35. *For each β there are only countably many $(\mathfrak{L}_b[L] \cup \mathfrak{L}_c[L])$ -structures up to isomorphism that are maximal at β .*

We can now put together these results to get the following.

Theorem 4.36 (Counterexample to Vaught's Conjecture). *The sentence $\sigma[L]$ has $\aleph_1$ many countable models up to isomorphism and has no perfect set of non-isomorphic countable models. In particular, $\sigma[L]$ is a counterexample to Vaught's conjecture for $\mathcal{L}_{\omega_1, \omega}$.*

Proof. We already have shown in Theorem 4.23 that $\mathbf{K}[L]$ is a weak Knight system, so by Theorem 2.25 it suffices to show Theorem 2.23 (a) – (d). First note that Theorem 2.23 (b) is exactly Theorem 4.35. Theorem 2.23 (d) follows immediately from Theorem 4.33.

Claim 4.37. *For every $\beta \in \omega_1$ and $\mathcal{M} \in K_\beta$ there is a unique $\mathcal{N} \in K$ such that $\tau_\beta(\mathcal{N}) = \mathcal{M}$ and $\mathcal{N}$ is maximal.*

Proof. For each $\mathcal{M} \in K_{-\infty}$ and $\beta \in \omega_1$ define $\mathcal{M}_\beta$ by induction as follows.

- Let $\mathcal{M}_0 \in K_0$ be such that $\mathcal{M}_0 \upharpoonright_{\mathfrak{L}_c[L]} = \mathcal{M}$. Note that such a $\mathcal{M}_0$ exists by the definition of a weak Knight system and is unique by Theorem 4.34 (Reconstruction).
- If β is a non-zero limit ordinal and $\mathcal{M}_i$ has been defined for each $i < \beta$, then by Theorem 4.34 (Limits) there must be a $\mathcal{M}_\beta \in K_\beta$ such that $\mathcal{M}_\beta \upharpoonright_{\mathfrak{L}_c[L]} = \mathcal{M}$. But by Theorem 4.34 (Reconstruction) such an $\mathcal{M}_\beta$ must be unique.
- If $\beta = \alpha + 1$ then let $\mathcal{M}_\beta$ be any element of K_β such that $\tau_\alpha(\mathcal{M}_\beta) = \mathcal{M}_\alpha$, if one exists. If one exists, by Theorem 4.34 (Reconstruction) it must be unique. If none exist stop the induction.

If this process stops, it must stop at some successor stage $\alpha + 1$. In this case, $\mathcal{M}_\alpha$ is the unique maximal lift of $\mathcal{M}$ and every lift of $\mathcal{M}$ is of the form $\tau_\zeta(\mathcal{M}_\alpha)$ for some $\zeta \leq \alpha$.

Suppose this process does not stop. By Theorem 4.34 (Boundedness), there is a $\gamma \in \omega_1$ such that for all $\zeta \geq \gamma$, we have $\text{rank}_\tau(\mathcal{M}_\zeta) \leq \gamma$, and so $\tau_\gamma(\mathcal{M}_\zeta) = \mathcal{M}_\zeta$. But then by Theorem 4.34 (Reconstruction), we must have $\mathcal{M}_\gamma = \tau_\gamma(\mathcal{M}_\zeta) = \mathcal{M}_\zeta$. Therefore $\mathcal{M}_\gamma$ is the unique maximal lift of $\mathcal{M}$ and every lift of $\mathcal{M}$ is of the form $\tau_\zeta(\mathcal{M}_\gamma)$ for some $\zeta \leq \gamma$. The claim follows. $\square$

Note that Theorem 2.23 (a) is immediate from Theorem 4.37. We also have from Theorem 4.34 (Reconstruction) and Theorem 4.37 that for each $\mathcal{M} \in K_{-\infty}$ there is a unique maximal lift and all other lifts are projections of it. Therefore $\mathcal{M}$ is eventually determined and Theorem 2.23 (c) holds. $\square$

REFERENCES

- [Ack15] N. L. Ackerman, *Vaught's conjecture without equality*, Notre Dame J. Formal Logic **56** (2015), no. 4, 573–582.
- [AFK26a] N. Ackerman, C. Freer, and R. Knight, *A counterexample to Vaught's conjecture for infinitary logic*, Palomar Registry, PALOMAR-2026-10-07-000001, version 1, October 2026, <https://palomar-registry.org/entry?id=PALOMAR-2026-10-07-000001&version=1>.
- [AFK26b] ———, *Vaught's conjecture counterexample: Lean 4 proof development*, GitHub repository, 2026, <https://github.com/cameronfreer/vaught-conjecture-counterexample/>.
- [Bec94] H. Becker, *The topological Vaught's conjecture and minimal counterexamples*, J. Symbolic Logic **59** (1994), no. 3, 757–784.
- [Bec13] ———, *Isomorphism of computable structures and Vaught's conjecture*, J. Symbolic Logic **78** (2013), no. 4, 1328–1344.
- [BFKL16] J. T. Baldwin, S. D. Friedman, M. Koerwien, and M. C. Laskowski, *Three red herrings around Vaught's conjecture*, Trans. Amer. Math. Soc. **368** (2016), no. 5, 3673–3694.
- [Bue08] S. Buechler, *Vaught's conjecture for superstable theories of finite rank*, Ann. Pure Appl. Logic **155** (2008), no. 3, 135–172.
- [Dre25] O. Dressler, *Lean theorem prover MCP*, <https://github.com/o0o0o0o/lean-lsp-mcp>, March 2025.
- [Fre25] C. E. Freer, *Lean 4 Skills: Theorem proving skill and workflow pack for AI coding agents*, <https://github.com/cameronfreer/lean4-skills>, October 2025.
- [Fre26] ———, *Infinitary logic in Lean 4*, <https://github.com/cameronfreer/infinitary-logic>, 2026.
- [Hjo01] G. Hjorth, *Vaught's conjecture on analytic sets*, J. Amer. Math. Soc. **14** (2001), no. 1, 125–143.
- [Hjo07] ———, *A note on counterexamples to the Vaught conjecture*, Notre Dame J. Formal Logic **48** (2007), no. 1, 49–51.
- [HS99] G. Hjorth and S. Solecki, *Vaught's conjecture and the Glimm-Effros property for Polish transformation groups*, Trans. Amer. Math. Soc. **351** (1999), no. 7, 2623–2641.
- [HSV94] B. Hart, S. Starchenko, and M. Valeriote, *Vaught's conjecture for varieties*, Trans. Amer. Math. Soc. **342** (1994), no. 1, 173–196.
- [Ive91] P. Iversen, *The number of countable isomorphism types of complete extensions of the theory of Boolean algebras*, Colloq. Math. **62** (1991), no. 2, 181–187.
- [Jec03] T. Jech, *Set theory*, third millenium ed., Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003.
- [Kni02] R. W. Knight, *The Vaught conjecture: A counterexample*, Unpublished manuscript, 2002.
- [Kur21] M. S. Kurilić, *Vaught's conjecture for almost chainable theories*, J. Symbolic Logic **86** (2021), no. 3, 991–1005.
- [LE65] E. G. K. Lopez-Escobar, *An interpolation theorem for denumerably long formulas*, Fund. Math. **57** (1965), 253–272.
- [May88] L. L. Mayer, *Vaught's conjecture for o-minimal theories*, J. Symbolic Logic **53** (1988), no. 1, 146–159.
- [Mon13] A. Montalbán, *A computability theoretic equivalent to Vaught's conjecture*, Adv. Math. **235** (2013), 56–73.
- [Mor70] M. Morley, *The number of countable models*, J. Symbolic Logic **35** (1970), 14–18.
- [Pun00] V. Puninskaya, *Vaught's conjecture for modules over a serial ring*, J. Symbolic Logic **65** (2000), no. 1, 155–163.
- [Sac83] G. E. Sacks, *On the number of countable models*, Southeast Asian conference on logic (Singapore, 1981), Stud. Logic Found. Math., vol. 111, North-Holland, Amsterdam, 1983, pp. 185–195.
- [Sam94] R. L. Sami, *Polish group actions and the Vaught conjecture*, Trans. Amer. Math. Soc. **341** (1994), no. 1, 335–353.
- [SHM84] S. Shelah, L. Harrington, and M. Makkai, *A proof of Vaught's conjecture for ω -stable theories*, Israel J. Math. **49** (1984), no. 1-3, 259–280.
- [Ste78] J. R. Steel, *On Vaught's conjecture*, Cabal Seminar 76–77 (Proc. Caltech-UCLA Logic Sem., 1976–77), Lecture Notes in Math., vol. 689, Springer, Berlin, 1978, pp. 193–208.
- [Ste21] ———, *On Vaught's conjecture*, Large cardinals, determinacy and other topics. The Cabal Seminar. Vol. IV, Lect. Notes Log., vol. 49, Cambridge Univ. Press, Cambridge, 2021, pp. 179–195.
- [Tan24] P. Tanović, *Vaught's conjecture for theories of discretely ordered structures*, Notre Dame J. Formal Logic **65** (2024), no. 3, 247–257.
- [Vau61] R. L. Vaught, *Denumerable models of complete theories*, Infinitistic Methods (Proc. Sympos. Foundations of Math., Warsaw, 1959), Pergamon, Oxford-London-New York-Paris, 1961, pp. 303–321.

HARVARD UNIVERSITY, CAMBRIDGE, MA, USA

Email address: nate@aleph0.net

MASSACHUSETTS INSTITUTE OF TECHNOLOGY, CAMBRIDGE, MA, USA

Email address: freer@mit.edu

UNIVERSITY OF OXFORD, OXFORD, UK

Email address: knight@maths.ox.ac.uk