Coherent extension and the countable classical excess

José Iovino

Department of Mathematics, The University of Texas at San Antonio, San Antonio, TX 78249, USA

We characterize simulation of a general probabilistic system of countable dimension by the countable classical target of Catani, Galley and Gonda: a simulation exists exactly when finite partial simulations can be extended coherently, without revising the masses and responses already assigned to finitely many ontic states. A transfinite derivative of these partial simulations ranks the failures of countable simulation; its first level recovers the countably additive excess of our earlier paper, and its fixed point decides, at each error, whether a countable simulation exists. On Cantor space we exhibit a system with countable simulations of every positive error and no exact one, and show that two demands defeat every exact simulation through a finite prepare-and-measure test chosen after the model has committed itself. We translate this defeat into a winning strategy in the refutation game of Aubone and Iovino for real-valued infinitary logic.

1 Introduction

In this paper we characterize simulation by the countable classical system $\Delta_{\mathbb{N}}$ of Catani, Galley and Gonda [1] through the coherent extension of partial simulations. The classical excess of a system A is the infimum $\varepsilon(A, \Delta_{\mathbb{N}})$ of the errors of the univalent simulations of A by $\Delta_{\mathbb{N}}$, that is, of the pairs of mixture-preserving maps that send preparations to probability vectors on $\mathbb{N}$ and effects to response functions on $\mathbb{N}$, the zero effect to the zero function, with the error measured uniformly over all preparations and effects.

The preceding paper [2] settles the corresponding question for countably additive classical targets. Its uniformization theorem states that a system has a countably additive classical simulation at a prescribed error exactly when each of its finitely generated subsystems has a finite classical simulation at that error; consequently the countably additive excess $\varepsilon_\sigma(A)$ is attained and is the supremum of the classical excesses of the finitely generated subsystems [2, Theorem A and Corollary 4.4]. The proof is by compactness, through a logical limit of finite models and Loeb's construction. The finite simulations it starts from are chosen independently of one another, and the logical limit imposes no compatibility among them.

For the countable target the uniformization fails, in two independent ways [2, Theorem B]. On Cantor space there is a system A whose finitely generated subsystems all have classical excess 0 while $\varepsilon(A, \Delta_{\mathbb{N}}) = 1$, and a system A', with the same states, whose classical excess is 0 and is not attained. In both, every stage is simulated exactly by a finite classical system. In the logical formulation of [2, Section 6], the conditions that an exact simulation of A' by $\Delta_{\mathbb{N}}$ would satisfy form a type, finitely approximately satisfiable in the

José Iovino: jose.iovino@utsa.edu, The most recent version of this paper is available at https://utsa-quantum-math.github.io/resources/coherent-extension-countable-excess/.

standard dual pair of Dueñez and omitted by it [2, Corollary 6.3]; this paper explains how it is omitted. What the countable target demands beyond the stages is *coherent extension*. The ontic states of a simulation by $\Delta_{\mathbb{N}}$ are discrete, and the restriction of such a simulation to finitely many ontic states and a finite stage is a partial simulation, which may retain less than the whole mass of each preparation, and which every later stage must extend: as the stages grow, the masses and responses already assigned to finitely many ontic states stay fixed, and the mass of each preparation not yet assigned must become small.

The second Cantor system shows why this matters to a physicist. Its preparations are the finite convex combinations of the conditional laws of the fair coin on the cylinders of Cantor space, the fair coin among them, and its effects are the cylinder-constant functions that are 1-Lipschitz for the natural ultrametric. Its classical excess is 0, approached by countable simulations of every positive error, and yet no countable simulation is exact. We show that the failure is detected in two rounds. A first demand compels any prospective exact model to place most of the mass of the fair coin on finitely many ontic states. Once it has done so, a second demand, the decomposition of the fair coin into its conditional laws on the cylinders of a depth chosen after the commitment, tested by effects scaled to that depth, cannot be met (Theorem B). The numerical excess cannot see this failure at all; the derivative introduced below sees it at its second level.

We make the requirement of coherent extension exact. Fix a system A of countable dimension, that is, an increasing union of bounded finite-dimensional systems, its *stages*. A *partial simulation* assigns to finitely many ontic states masses for the states of a stage and responses for its effects, subject to two inequalities that a completion at error r requires; a *demand* asks for a later stage and a smaller unassigned mass; and a *consistency property* at error r is a nonempty family of partial simulations in which every member has, for every demand, an extension in the family that meets it (Section 4). Removing, transfinitely, the partial simulations that cannot meet every demand by an extension in the family gives decreasing families $\mathcal{C}_\alpha(r)$ of partial simulations, indexed by ordinals, and the *excess of level* α, $\varepsilon^{(\alpha)}(A)$, is the infimum of the errors r at which the empty partial simulation survives to $\mathcal{C}_\alpha(r)$ (Section 5).

Theorem A (Coherent extension)**.** *Let A be a system of countable dimension, with a fixed exhaustion by stages, and let $r \geq 0$.*

(i) *A has a univalent simulation by $\Delta_{\mathbb{N}}$ with error at most r if and only if A has a consistency property at error r, if and only if the families $\mathcal{C}_\alpha(r)$ stabilize at a nonempty family; and a partial simulation belongs to that family if and only if it is the restriction of such a simulation.*

(ii) *The empty partial simulation belongs to $\mathcal{C}_1(r)$ if and only if $r \geq \varepsilon_\sigma(A)$; in particular $\varepsilon^{(1)}(A) = \varepsilon_\sigma(A)$.*

(iii) $\varepsilon(A, \Delta_{\mathbb{N}}) = \sup_\alpha \varepsilon^{(\alpha)}(A)$.

(iv) *For two exhaustions each of whose stages is contained in a stage of the other, the excesses of every level, and the level at which the empty partial simulation is removed at error r, coincide.*

Parts (i)–(iii) are Proposition 4.3 and Theorem 5.4, and part (iv) is Proposition 5.5. By part (ii), the uniformization theorem of [2] is the first level of this analysis: compactness decides exactly the obstructions that a single stage presents by itself, and the higher levels record those that appear only when finite simulations must cohere. The classical excess is finitely determined, that is, equal to $\varepsilon_\sigma(A)$, exactly when no level adds to the first.

The levels have a direct operational reading, as a contest between a modeler, who commits masses and responses on finitely many ontic states, and a challenger, who demands more of the system without allowing revision; we describe it in Section 3.

The central example is the Lipschitz Cantor system A', which separates the two invariants that the derivative provides; the nested Cantor system A is defeated in the same way, at every error below 1.

Theorem B (Two demands defeat exact simulation)**.** *For every exhaustion of the Cantor systems by stages, the following hold.*

(1) *For the Lipschitz Cantor system A', $\varepsilon^{(\alpha)}(A') = 0$ for every ordinal α, while at error 0 the empty partial simulation belongs to $\mathcal{C}_1(0)$ and not to $\mathcal{C}_2(0)$.*

(2) *For the nested Cantor system A, $\varepsilon^{(1)}(A) = 0$ and $\varepsilon^{(2)}(A) = \varepsilon(A, \Delta_{\mathbb{N}}) = 1$.*

In both systems the defeat is the localization of Lemma 6.1, which also gives two lower bounds for the second level of an arbitrary system, by the *fine-scale distinguishability* $D(A)$ of its preparations and by their *fine-scale separation* $\mathrm{Sep}(A)$ (Proposition 7.2):

$$\varepsilon^{(2)}(A) \geq \max\{\varepsilon_\sigma(A), D(A), \tfrac{1}{2}\,\mathrm{Sep}(A)\}.$$

Neither bound dominates the other, and stability of the pairing, in the sense of Grothendieck's double limit condition, excludes neither (Proposition 7.3); the bound can be strict, even for a stable system with an exact Borel model (Proposition 7.4). In every system we have computed, the classical excess is determined at the second level.

Consistency properties are the device that takes the place of compactness in infinitary logic, where countable conjunctions and disjunctions express, among other things, countable additivity [3, 4]. For real-valued infinitary logic, the model existence theorem is due to Aubone, Dueñez and Iovino [5], and Aubone and Iovino organize it as a game whose refutations are well-founded trees of finite contradictions [6]. The derivative of this paper is stated directly in terms of the system and uses no syntax; in Section 8 we write the countable target as a theory in their logic and translate the defeat of exact simulation of the Lipschitz Cantor system into a winning strategy in their game, in which the answer to a single countable challenge selects the finite experiment, and the later answers only fix how far its verification must reach (Proposition 8.2).

Organization of the paper. Section 2 recalls the setting of [2] and the Cantor systems. Section 3 describes the challenge procedure informally, Section 4 defines partial simulations and consistency properties, and Section 5 defines the derivative and proves Theorem A. Section 6 proves Theorem B, and Section 7 treats the second level in general, with stable examples. Section 8 writes the countable target as a theory of real-valued infinitary logic and translates the defeat of exact simulation into the refutation game of Aubone and Iovino. Sections 9 and 10 give two complements: analytic conditions under which the first level suffices, and attainment at the countably additive optimum. Section 11 collects open questions. The appendix constructs an optimal finite simulation for one of the stable examples.

2 Systems and the countable target

We use the definitions of [2, Section 2]. A *system* A consists of a nonempty convex set Ω_A of states, a convex set E_A of effects containing 0, and a pairing $(e, \omega) \mapsto e \cdot \omega \in [0, 1]$,

affine in each argument, with $0 \cdot \omega = 0$. For a nonempty measurable space $(\Omega, \mathcal{F})$, the classical system $\Delta_{(\Omega,\mathcal{F})}$ has the countably additive probability measures as states and the measurable functions $\Omega \to [0,1]$ as effects, with pairing $f \cdot \gamma = \int f \, d\gamma$, the simulated probability of a response against a preparation measure; for a finite or countable set we take the power set and write $\Delta_{\mathbb{N}}$ and $\Delta_k = \Delta_{\{1,\dots,k\}}$. A *univalent ϵ-simulation* of A by a classical system is a pair of mixture-preserving maps Γ on states and Θ on effects, with $\Theta(0) = 0$ and $|e \cdot \omega - \Theta(e) \cdot \Gamma(\omega)| \leq \epsilon$ for all states ω and effects e; its least such ϵ is its *error*. We write $\varepsilon(A, B)$ for the infimum of the errors of univalent simulations of A by a classical system B, and $\varepsilon_\sigma(A)$ for the infimum of $\varepsilon(A, \Delta_{(\Omega,\mathcal{F})})$ over all measurable spaces. By [2, Lemma 2.1], the map Θ of a univalent simulation extends to a linear map on the span of the effects, and in particular $\Theta(\sum_j e_j) = \sum_j \Theta(e_j)$ whenever $e_1, \dots, e_k$ and $\sum_j e_j$ are effects.

A system A has *countable dimension* if it is the increasing union of systems $A_1 \subseteq A_2 \subseteq \cdots$, its *stages*, whose states and effects form bounded subsets of finite-dimensional spaces; we write Ω_n and E_n for the states and effects of A_n, and call the sequence of stages an *exhaustion*. Every finitely generated subsystem of A lies in some stage. For such systems [2, Corollaries 4.4 and 4.5] give

$$\varepsilon_\sigma(A) = \sup_F \varepsilon(F, \Delta_{\mathbb{N}}) = \lim_n \varepsilon(A_n, \Delta_{\mathbb{N}}),$$

the supremum over the finitely generated subsystems F of A, and each stage has a simulation by a finite classical system whose error is $\varepsilon(A_n, \Delta_{\mathbb{N}})$ [2, Theorem 3.2]. We say that the classical excess of A is *finitely determined* if $\varepsilon(A, \Delta_{\mathbb{N}}) = \varepsilon_\sigma(A)$.

The Cantor systems. Let $K = \{0,1\}^{\mathbb{N}}$, with the fair-coin probability p and the ultrametric ρ for which two distinct sequences whose longest common initial segment has length j are at distance 2^{-j}. For an integer $n \geq 0$ let $\mathcal{P}_n$ be the partition of K into the 2^n cylinders of depth n, and for $C \in \mathcal{P}_n$ let p_C be the conditional law of p on C. Let A_n be the system whose states are the convex combinations of the laws p_C with $C \in \mathcal{P}_n$, whose effects are the functions $K \to [0,1]$ constant on each member of $\mathcal{P}_n$, and whose pairing is integration, and let A be the increasing union of the systems A_n. Let A'_n have the states of A_n and the effects of A_n that are 1-Lipschitz for ρ, and let A' be the increasing union of the systems A'_n. These are the systems of [2, Theorem B]: $\varepsilon_\sigma(A) = \varepsilon_\sigma(A') = 0$, $\varepsilon(A, \Delta_{\mathbb{N}}) = 1$, and $\varepsilon(A', \Delta_{\mathbb{N}}) = 0$, while no simulation of A' by $\Delta_{\mathbb{N}}$ has error 0. Both have countable dimension, with the exhaustions by depth, but Theorem B holds for every exhaustion.

3 The challenge procedure

A modeler proposes to simulate A by $\Delta_{\mathbb{N}}$ at error r, and a challenger tests the proposal. At each moment the modeler has committed to finitely many ontic states, to the masses that the states of some stage place on them, and to the responses that the effects of that stage give at them. The commitments must leave room for a completion at error r: for every state and effect of the stage, the probability already accounted for may not exceed the true probability by more than r, and may not fall short of it by more than r plus the mass not yet assigned.

The challenger then makes a *demand*: a later stage and a bound δ on the mass that any state of that stage may leave unassigned. The modeler must answer with new commitments that meet the demand and *extend* the old ones, adding ontic states and assigning masses and responses for the larger stage, without changing anything already committed. A

countable simulation at error r is exactly a strategy that answers every demand forever (Proposition 4.3).

The rounds measure how far a commitment is from defeat. A commitment survives one round if every demand can be answered by an admissible commitment that extends it; by Theorem 5.4, the empty commitment survives one round exactly when $r \geq \varepsilon_\sigma(A)$, which is the content of the uniformization theorem of [2]. A commitment survives two rounds if every demand can be answered by a commitment that itself survives one round, and so on; at transfinite levels the challenger announces, before each demand, an ordinal smaller than the last one announced, and the modeler survives to level α if every such countdown from α can be answered (Remark 5.3).

Consider the Lipschitz Cantor system at error 0. The first demand leaves at most a quarter of the mass of the fair coin unassigned, and the modeler answers by committing a finite set H of ontic states. The challenger then chooses a depth n with $2^n > 4|H|$. The second demand includes the conditional laws of the fair coin on the cylinders of depth n and the effects $2^{-n}1_C$, and bounds the remaining mass by a quarter of the common scale 2^{-n} of those effects; by Theorem B, no extension meets it. Every finite subsystem has an exact finite simulation, so that no finite test chosen in advance excludes an exact countable simulation; the test that does is chosen after the commitment.

The procedure also describes a method that applies beyond the examples of this paper: compel a prospective ontological model to commit, and then find the finite prepare-and-measure test that defeats the commitment.

4 Partial simulations and consistency properties

Fix a system A of countable dimension, an exhaustion $A_1 \subseteq A_2 \subseteq \cdots$, and a real number $r \geq 0$.

Definition 4.1. A *partial simulation* of A at error r, of *level* n, consists of a finite set $H \subseteq \mathbb{N}$ of ontic states, an affine map γ from Ω_n to the subprobability vectors on H, and an affine map θ from E_n to $[0,1]^H$ with $\theta(0) = 0$, such that, with $d(\omega) = 1 - \sum_{x \in H} \gamma(\omega)(x)$,

$$\theta(e) \cdot \gamma(\omega) \leq e \cdot \omega + r \qquad \text{and} \qquad e \cdot \omega \leq \theta(e) \cdot \gamma(\omega) + r + d(\omega)$$

for all $\omega \in \Omega_n$ and $e \in E_n$. Its *defect* is $\sup_{\omega \in \Omega_n} d(\omega)$. A partial simulation (H', γ', θ') of level n' *extends* (H, γ, θ) if $n' \geq n$, $H \subseteq H'$, and $\gamma'(\omega)$ and $\theta'(e)$ restrict on H to $\gamma(\omega)$ and $\theta(e)$ for all $\omega \in \Omega_n$ and $e \in E_n$. A *demand* is a pair (n, δ) of a positive integer and a real $\delta > 0$, and a partial simulation *meets* it if its level is at least n and its defect is at most δ.

The number $d(\omega)$ is the mass of ω not yet assigned. In any completion, the simulated probability of e at ω lies between the assigned contribution $\theta(e) \cdot \gamma(\omega)$ and that contribution plus $d(\omega)$. The two inequalities require this interval to meet the interval of radius r about $e \cdot \omega$; they give a necessary condition for each preparation and effect separately, and coherent extension asks whether one assignment can meet the requirements together. The *empty* partial simulation, with $H = \varnothing$ and level 1, satisfies both inequalities, and every partial simulation extends it. As for simulations, the argument of [2, Lemma 2.1] shows that $\theta(\sum_j e_j) = \sum_j \theta(e_j)$ whenever $e_1, \ldots, e_k$ and $\sum_j e_j$ belong to E_n.

Definition 4.2. A *consistency property* for A at error r is a nonempty family $\mathcal{C}$ of partial simulations at error r such that every member of $\mathcal{C}$ has, for every demand, an extension in $\mathcal{C}$ that meets the demand.

Proposition 4.3. *Let A be a system of countable dimension and let $r \geq 0$. Then A has a univalent simulation by $\Delta_{\mathbb{N}}$ with error at most r if and only if A has a consistency property at error r. Moreover, every member of a consistency property at error r is the restriction of such a simulation, and the restrictions of such a simulation form a consistency property. In particular, $\varepsilon(A, \Delta_{\mathbb{N}})$ is the infimum of the reals r at which A has a consistency property.*

Proof. Let (Γ, Θ) be a univalent simulation of A by $\Delta_{\mathbb{N}}$ with error at most r, and let $\mathcal{C}$ be the family of its restrictions to finite sets $H \subseteq \mathbb{N}$ and stages A_n. Restriction lowers a simulated probability by at most the mass removed, so each restriction satisfies the two inequalities. The states of a stage A_n are affine combinations, with bounded coefficients, of finitely many of its states, so that the masses of all states of A_n outside a large finite set H are uniformly small; hence every demand is met by an extension within $\mathcal{C}$.

Conversely, let $\mathcal{C}$ be a consistency property, let c_0 be a member of $\mathcal{C}$, and choose in $\mathcal{C}$ a sequence $c_1, c_2, \ldots$, each extending the one before, with c_k meeting the demand $(k, 2^{-k})$ for $k \geq 1$. The union of the finite sets H_k is a subset of $\mathbb{N}$; on it the maps γ_k and θ_k define mixture-preserving maps Γ on the states and Θ on the effects of A, extended by 0 off the union, with $\Theta(0) = 0$, and every $\Gamma(\omega)$ has mass 1 because the defects tend to 0. For ω and e in a stage A_n, the probabilities $\theta_k(e) \cdot \gamma_k(\omega)$ increase to $\Theta(e) \cdot \Gamma(\omega)$, and the two inequalities pass to the limit. The simulation so obtained restricts to c_0. □

The proof is a chain construction, and the content of the proposition is the reformulation: a simulation by $\Delta_{\mathbb{N}}$ is the union of a coherent chain of partial simulations, where the uniformization theorem of [2] produces a logical limit. Condition (i) of that theorem asks each finitely generated subsystem to be simulated by itself; a consistency property asks in addition that the finite simulations cohere, through extensions that never revise what has been assigned.

5 The derivative

Every partial simulation either belongs to a consistency property or does not, and in the second case the failure has a rank, computed by a derivative of the kind familiar from Cantor–Bendixson and Scott analysis.

Definition 5.1. Let $\mathcal{C}_0(r)$ be the family of all partial simulations of A at error r. For an ordinal α, let $\mathcal{C}_{\alpha+1}(r)$ consist of the members of $\mathcal{C}_\alpha(r)$ that have, for every demand, an extension in $\mathcal{C}_\alpha(r)$ meeting it, and for a limit ordinal λ let $\mathcal{C}_\lambda(r) = \bigcap_{\alpha<\lambda} \mathcal{C}_\alpha(r)$. The *excess of level* α of A is the infimum $\varepsilon^{(\alpha)}(A)$ of the reals r for which the empty partial simulation belongs to $\mathcal{C}_\alpha(r)$.

The families $\mathcal{C}_\alpha(r)$ decrease in α and increase in r, so that $\varepsilon^{(\alpha)}(A)$ is nondecreasing in α. They stabilize at a family $\mathcal{C}_\infty(r)$, which is the largest consistency property at error r when it is nonempty, since every consistency property lies in every $\mathcal{C}_\alpha(r)$. The *removal level* at error r of a partial simulation outside $\mathcal{C}_\infty(r)$ is the least ordinal α such that it lies outside $\mathcal{C}_\alpha(r)$.

Lemma 5.2. *If a partial simulation c' extends a partial simulation c and belongs to $\mathcal{C}_\alpha(r)$, then c belongs to $\mathcal{C}_\alpha(r)$.*

Proof. By induction on α. At a successor, an extension of c' in $\mathcal{C}_\alpha(r)$ that meets a demand is also an extension of c, and c belongs to $\mathcal{C}_\alpha(r)$ by induction; at limits one takes intersections. □

Remark 5.3. By induction on α, a partial simulation c belongs to $\mathcal{C}_\alpha(r)$ exactly when for every ordinal $\beta < \alpha$ and every demand there is an extension of c in $\mathcal{C}_\beta(r)$ that meets it. This is the countdown game of Section 3: the challenger announces $\beta_1 < \alpha$ and a demand, the modeler answers by an extension, the challenger announces $\beta_2 < \beta_1$ and a demand, and so on until the ordinal 0 is announced, after which no further demand is made; the partial simulation lies in $\mathcal{C}_\alpha(r)$ exactly when the modeler can always answer. For a positive integer k, membership in $\mathcal{C}_k(r)$ is the ability to answer k successive demands.

Theorem 5.4. *Let A be a system of countable dimension and let $r \geq 0$.*

(a) *The empty partial simulation belongs to $\mathcal{C}_1(r)$ if and only if $r \geq \varepsilon_\sigma(A)$; in particular $\varepsilon^{(1)}(A) = \varepsilon_\sigma(A)$.*

(b) *A has a univalent simulation by $\Delta_\mathbb{N}$ with error at most r if and only if $\mathcal{C}_\infty(r)$ is nonempty, and a partial simulation belongs to $\mathcal{C}_\infty(r)$ if and only if it is the restriction of such a simulation.*

(c) $\varepsilon(A, \Delta_\mathbb{N}) = \sup_\alpha \varepsilon^{(\alpha)}(A)$.

Proof. (a) The empty partial simulation belongs to $\mathcal{C}_1(r)$ exactly when every demand (n, δ) is met by some partial simulation at error r. Such a partial simulation, completed by one more ontic state that carries the unassigned mass $d(\omega)$ and responds to each effect e with $e \cdot q$, for a fixed state q of A_n, is a simulation of A_n by a finite classical system with error at most $r + \delta$; hence $\varepsilon(A_n, \Delta_\mathbb{N}) \leq r$ for every n, and $\varepsilon_\sigma(A) \leq r$. Conversely, if $r \geq \varepsilon_\sigma(A)$, each stage A_n has a simulation by a finite classical system with error $\varepsilon(A_n, \Delta_\mathbb{N}) \leq \varepsilon_\sigma(A) \leq r$, which meets the demand (n, δ) with defect 0.

(b) The family $\mathcal{C}_\infty(r)$, when nonempty, is a consistency property, and every consistency property is contained in it; both assertions therefore follow from Proposition 4.3.

(c) Each $\varepsilon^{(\alpha)}(A)$ is at most $\varepsilon(A, \Delta_\mathbb{N})$, since the restrictions of a simulation with error at most r form a consistency property containing the empty partial simulation. If $r < \varepsilon(A, \Delta_\mathbb{N})$, the empty partial simulation is not in $\mathcal{C}_\infty(r) = \mathcal{C}_\alpha(r)$ for some ordinal α, and by monotonicity in r it is not in $\mathcal{C}_\alpha(r')$ for any $r' \leq r$, so that $\varepsilon^{(\alpha)}(A) \geq r$. □

The families $\mathcal{C}_\alpha(r)$, the removal levels and the level excesses are defined relative to the chosen exhaustion. By Theorem 5.4, two invariants do not depend on it: $\varepsilon^{(1)}(A) = \varepsilon_\sigma(A)$, and the nonemptiness of $\mathcal{C}_\infty(r)$, which holds exactly when a simulation by $\Delta_\mathbb{N}$ with error at most r exists; the families themselves live on domains that depend on the stages. For exhaustions that are mutually cofinal by containment the intermediate levels agree as well.

Proposition 5.5. *Let (A_n) and (B_m) be exhaustions of A such that every stage of each is contained in a stage of the other, and let $r \geq 0$. For every ordinal α, the restriction of a member of $\mathcal{C}_\alpha(r)$ for (A_n) to a stage B_m contained in its domain is a member of $\mathcal{C}_\alpha(r)$ for (B_m), and symmetrically. The empty partial simulation has the same removal level at error r for both exhaustions, and the excesses of every level coincide.*

Proof. At level 0, restriction preserves the two inequalities of Definition 4.1 and does not increase the defect. At a successor $\alpha + 1$, let a member c of $\mathcal{C}_{\alpha+1}(r)$ for (A_n) be restricted to B_m, and let a demand for (B_m) ask for the stage B_k and defect at most δ. Choose a stage B_j containing B_k and B_m, and a stage A_l containing B_j. The member c has an extension in $\mathcal{C}_\alpha(r)$ that meets the demand (l, δ); its restriction to B_j lies in $\mathcal{C}_\alpha(r)$ for (B_m) by induction, meets the original demand, and extends the restriction of c. At limit levels one takes intersections. For the empty partial simulations, at a successor, a member of the preceding level for (A_n) on a stage containing both the demanded stage and B_1 restricts,

by the same comparison, to an extension of the empty partial simulation for (B_m) in the preceding level; the reverse argument gives equality. □

The hypothesis is cofinality by containment, which an arbitrary pair of exhaustions need not satisfy. Together, Theorem 5.4 and Proposition 5.5 prove Theorem A.

6 Two demands defeat exact simulation

The estimate behind Theorem B compares the unassigned mass with the scale of the effects at the selected depth: at scale 1 it gives an error gap, and at scale 2^{-n} it detects only the endpoint. Both Cantor systems are defeated by this one estimate, a form of the matching lemma of [2, Lemma 5.1] for partial simulations.

Lemma 6.1 (Localization)**.** *Let $r \geq 0$ and $\delta, \delta' > 0$, let p be a state of a stage A_{n_0}, and let (H, γ, θ) be a partial simulation at error r, of level at least n_0, such that $\gamma(p)$ has mass at least $1 - \delta$ on H. Let n be a positive integer such that the constant 1 is an effect of A_n, let $c > 0$, and suppose that $p = \frac{1}{N}\sum_C p_C$ is the uniform average of N states p_C of A_n for which the functions $c1_C$ are effects of A_n with sum $c \cdot 1$ and $c1_C \cdot p_C = c$. Then every extension of (H, γ, θ) that meets the demand (n, δ') satisfies*

$$c\Big(1 - \frac{|H|}{N} - \delta\Big) \leq r + \delta'.$$

Proof. Let (H', γ', θ') be such an extension, with unassigned masses d'. The second inequality of Definition 4.1, averaged over the decomposition, gives $c \leq \frac{1}{N}\sum_C \theta'(c1_C) \cdot \gamma'(p_C) + r + \delta'$, since the average of the numbers $d'(p_C)$ is $d'(p) \leq \delta'$. Additivity and $\theta'(0) = 0$ give $\sum_C \theta'(c1_C) = \theta'(c \cdot 1) = c\,\theta'(1) \leq c$ at every ontic state, so that the part of the average carried by H is at most $c|H|/N$; outside H each $\theta'(c1_C)$ is at most c, and $\gamma'(p)$ has mass at most δ there. □

Proof of Theorem B. Fix an exhaustion of A or A' by stages. The values $\varepsilon^{(1)} = \varepsilon_\sigma$ and $\varepsilon(\,\cdot\,, \Delta_{\mathbb{N}})$ are those of [2, Theorem B], and they bound every $\varepsilon^{(\alpha)}$ by Theorem 5.4. Let p be the fair coin, let n_0 be such that p is a state of the stage of index n_0, and for a cylinder C of depth n let p_C be the conditional law of p on C, with $N = 2^n$. If the empty partial simulation belongs to $\mathcal{C}_2(r)$, then for each $\delta > 0$ some (H, γ, θ) in $\mathcal{C}_1(r)$ meets the demand (n_0, δ), so that $\gamma(p)$ has mass at least $1 - \delta$ on H, and every demand is met by one of its extensions. Once H is fixed, the decomposition of p into the laws p_C of depth n involves finitely many states and effects, which lie in some stage of the exhaustion, of index m, together with the constant 1; we apply Lemma 6.1 with the demand (m, δ'). For part (1), let $r = 0$ and take $c = 2^{-n}$, for which the functions $c1_C$ are 1-Lipschitz for ρ and therefore effects of A', since a point of C and a point outside C are at distance at least $2^{-(n-1)}$; take $\delta < \frac{1}{4}$, n with $|H|/N < \delta$, and $\delta' < c/4$. The lemma gives $(1 - 2\delta)c < \delta'$, whose left side exceeds $c/2$ while its right side is below $c/4$. Moreover, ε_σ of every stage of A' is 0, since the identity model on Cantor space is exact, and every stage, being finite-dimensional, has a simulation by a finite classical system with error 0 [2, Theorem 3.2]; so every demand is met at error 0 and the empty partial simulation belongs to $\mathcal{C}_1(0)$. For part (2), take $c = 1$: the lemma gives $1 - |H|/N - \delta \leq r + \delta'$, and letting n grow and δ, δ' shrink gives $r \geq 1$. □

In both systems the second demand refines what the first demand localized: the first fixes a finite set H of ontic states carrying most of the mass of the fair coin, and a decomposition into more cylinders than H can resolve defeats every extension. For A' the defeat

occurs at error 0 only: it is the non-attainment of [2, Theorem B(2)], located at the second level, in a system whose every numerical level of excess is 0. For A it occurs at every error below 1.

7 The second level

The localization behind Theorem B applies to every system and gives two lower bounds for the second level.

Definition 7.1. Let p be a state of a system A. For $w > 0$, let $D_w(p)$ be the supremum of $\sum_j w_j\,(e_j \cdot p_j)$ over all decompositions $p = \sum_{j=1}^{k} w_j p_j$ into states p_j of A, not necessarily distinct, with weights $0 < w_j \le w$ and $\sum_j w_j = 1$, and all effects $e_1, \dots, e_k$ of A whose sum is an effect of A. The *fine-scale distinguishability* of p is $D(p) = \inf_{w>0} D_w(p)$. For a real $s > 0$, a decomposition $p = \frac{1}{N}\sum_{j=1}^{N} p_j$ is *s-separated* if for all $j \ne k$ some effect e of A satisfies $|e \cdot p_j - e \cdot p_k| \ge s$, and the *fine-scale separation* $\mathrm{Sep}(p)$ is the supremum of the reals s for which p has s-separated decompositions with arbitrarily large N, or 0 if there are none. Put $D(A) = \sup_p D(p)$ and $\mathrm{Sep}(A) = \sup_p \mathrm{Sep}(p)$.

Proposition 7.2. *For every system A of countable dimension,*

$$\varepsilon^{(2)}(A) \ge \max\{\varepsilon_\sigma(A), D(A), \tfrac{1}{2}\,\mathrm{Sep}(A)\}.$$

Proof. The bound by $\varepsilon_\sigma(A) = \varepsilon^{(1)}(A)$ is monotonicity. Let p be a state of a stage A_{n_0}, let $r \ge 0$, and let $0 < \delta < 1$. If the empty partial simulation lies in $\mathcal{C}_2(r)$, some $c = (H, \gamma, \theta)$ in $\mathcal{C}_1(r)$ meets the demand (n_0, δ), so that $\gamma(p)$ has mass at least $1 - \delta > 0$ on H, which is therefore nonempty, and every demand (n, δ') is met by some extension $c' = (H', \gamma', \theta')$ of c.

Let $p = \sum_j w_j p_j$ be a decomposition, with effects e_j whose sum is an effect, choose a stage A_n containing the states p_j, the effects e_j and their sum, and let c' meet (n, δ'). Summing the second inequality of Definition 4.1 with the weights w_j, and using that $\sum_j w_j d'(p_j) = d'(p) \le \delta'$, that $\sum_j \theta'(e_j) = \theta'(\sum_j e_j) \le 1$ at every ontic state, and that $\gamma'(p)$ has mass at most δ outside H, gives

$$\sum_j w_j\,(e_j \cdot p_j) \le |H| \max_j w_j + \delta + r + \delta'.$$

Taking w small with $|H|w \le \delta$ gives $D(p) \le r + 2\delta + \delta'$.

For the separation, let $p = \frac{1}{N}\sum_j p_j$ be s-separated, choose for each pair $j \ne k$ a separating effect, choose a stage A_n containing the states p_j and these finitely many effects, and let c' meet (n, δ'). Fewer than $N/2$ of the vectors $\gamma'(p_j)$ have mass above 2δ outside H. If N exceeds twice the number M of cubes of side $\delta/|H|$ needed to cover $[0,1]^H$, two of the others, $\gamma'(p_j)$ and $\gamma'(p_k)$ with $j \ne k$, satisfy $\|\gamma'(p_j) - \gamma'(p_k)\|_1 \le 5\delta$. For an effect e that separates p_j from p_k, the two inequalities of Definition 4.1 give $s \le 2r + 5\delta + \delta'$; since M depends only on $|H|$ and δ, such N exist whenever $s < \mathrm{Sep}(p)$.

Since δ and δ' are arbitrary, the empty partial simulation lies outside $\mathcal{C}_2(r)$ whenever $r < D(p)$ or $r < \frac{1}{2}\,\mathrm{Sep}(p)$. □

The two bounds are independent, and stability of the pairing excludes neither. Let $B_1, B_2, \dots$ be disjoint sets of coordinates of K with $|B_n| = n$, and let $\mathcal{Q}_n$ be the partition of K into the 2^n cells determined by the coordinates in B_n, with $\mathcal{Q}_0 = \{K\}$; for a cell C let p_C be the conditional law of p on C. Let S be the system whose states are the convex

combinations of the laws p_C of all cells, whose effects are the convex combinations of 0 and the indicators 1_C, and whose pairing is integration; for a real $c \in (0,1)$ let S_c have the same states and the effects ce with e an effect of S; and, to suppress the matching score of effects with admissible sum while preserving pairwise separation, for a real $b \in (0,1)$ let T have the same states and the effects $\kappa(b + (1-b)e)$, with $\kappa \in [0,1]$ and e an effect of S. The pairing of a system is *stable* if $\lim_n \lim_m e_n \cdot \omega_m = \lim_m \lim_n e_n \cdot \omega_m$ for all sequences of effects e_n and states ω_m for which both iterated limits exist; this is Grothendieck's double limit condition and the stability, in the sense of model theory, of the pairing. The system A, built on nested cylinders, has the order property instead: for nested cylinders $C_1 \supseteq C_2 \supseteq \cdots$ of depths $1, 2, \ldots$, the two iterated limits of $1_{C_n} \cdot p_{C_m}$ are 1 and 0.

Proposition 7.3. *The systems S, S_c and T have countable dimension and stable pairings, and $\varepsilon_\sigma = 0$ for each of them. Moreover*

$$\varepsilon(S, \Delta_{\mathbb{N}}) = \varepsilon^{(2)}(S) = D(S) = 1, \qquad \varepsilon(S_c, \Delta_{\mathbb{N}}) = \varepsilon^{(2)}(S_c) = D(S_c) = c,$$

while $D(T) = 0$ and $\varepsilon(T, \Delta_{\mathbb{N}}) = \varepsilon^{(2)}(T) = \frac{1}{2}\operatorname{Sep}(T) = \frac{1}{2}(1-b)$.

Proof. Finite parts. The densities with respect to p of the states and the effects that involve only the blocks $B_1, \ldots, B_N$ are measurable for the finite algebra those coordinates generate, and recording masses and values on its atoms simulates them exactly; hence $\varepsilon_\sigma = 0$ for all three systems.

Stability. For cells $C \in \mathcal{Q}_n$ and $C' \in \mathcal{Q}_m$, independence of the blocks gives $1_C \cdot p_{C'} = 2^{-n}$ if $n \neq m$, and $1_C \cdot p_{C'} = 1$ or 0 according as $C = C'$ or not, if $n = m$. Write an effect of S as $\sum_g \mu_g 1_g$ and a state as $\sum_h \lambda_h p_h$, with μ a subprobability vector and λ a probability vector on the countable set of cells, and let $\mu(n)$ and $\lambda(n)$ be their masses on level n; then the pairing is

$$P(\mu, \lambda) = \sum_g \mu_g \lambda_g + \sum_n 2^{-n} \mu(n) - \sum_n 2^{-n} \mu(n) \lambda(n).$$

Passing to subsequences of representing vectors that converge pointwise, the level masses converge, the weights 2^{-n} are summable, and the first term converges by dominated convergence against whichever vector is fixed, so that both iterated limits equal $P(\mu, \lambda)$; for S_c and T one passes in addition to a subsequence along which the scalar factors converge.

Values. The fair coin is the uniform average of the laws p_C with $C \in \mathcal{Q}_n$, the effects 1_C sum to 1, and $1_C \cdot p_C = 1$; this gives $D(S) = 1$ and $D(S_c) \geq c$, while every effect of S_c takes values at most c, so that $D(S_c) \leq c$ and the zero response has error at most c. For T, the effects $b + (1-b)1_C$ separate the laws p_C of one level by $1-b$, and no effect of T varies by more than $1-b$. If effects $\kappa_j(b + (1-b)e_j)$ of T have a sum that is an effect, then $\sum_j \kappa_j b \leq 1$, so that $D_w(q) \leq w/b$ for every state q and $D(T) = 0$. Proposition 7.2 gives the lower bounds, and the classical excesses bound $\varepsilon^{(2)}$ from above by Theorem 5.4. For T, a simulation by a finite classical system with error $\frac{1}{2}(1-b)$ is constructed in Appendix A. □

A single preparation can thus obstruct countable simulation in two distinct ways at the second level: by decompositions whose members are distinguished by effects whose sum is an effect, and by decompositions whose members are separated in pairs. Neither bound dominates the other, since $\frac{1}{2}\operatorname{Sep}(S_c) \leq c/2 < D(S_c)$ and $D(T) = 0 < \frac{1}{2}\operatorname{Sep}(T)$.

The bound of Proposition 7.2 can be strict. The fine-scale distinguishability D weighs each distinguishing effect by its size, whereas small effects carry positivity constraints that survive division by their scale: a family of effects at a scale δ_n tending to 0 contributes

nothing to a matching score in the limit, yet constrains the responses to the other effects enough to carry the full localization argument. For $b \in (0,1)$ and reals $0 < \delta_n \le 1-b$ with $\delta_n \to 0$, let T^+ be the system with the states of S whose effects are the convex combinations of 0, of the functions $b + (1-b)e$ with e an effect of S, and of the functions $\delta_n 1_C$ with $C \in \mathcal{Q}_n$ and $n \ge 1$.

Proposition 7.4. *The system T^+ has countable dimension and a stable pairing, $\varepsilon_\sigma(T^+) = 0$, $D(T^+) = 0$, $\mathrm{Sep}(T^+) = 1-b$, and*

$$\varepsilon(T^+, \Delta_{\mathbb{N}}) = \varepsilon^{(2)}(T^+) = 1-b > \tfrac{1}{2}(1-b) = \max\{\varepsilon_\sigma(T^+), D(T^+), \tfrac{1}{2}\,\mathrm{Sep}(T^+)\}.$$

Proof. The identity model on Cantor space is exact, the finite parts are simulated exactly as for S, and stability follows from the calculation for S after the bounded constant part of each effect is separated from its combination of cell indicators. No generating effect varies over the states by more than $1-b$, so that $\mathrm{Sep}(T^+) \le 1-b$, and the effects $b + (1-b)1_C$ separate the laws p_C of one level by $1-b$. Sending every state to one point mass and every effect f to the constant $f \cdot p$ gives a simulation with error at most $1-b$, so that $\varepsilon(T^+, \Delta_{\mathbb{N}}) \le 1-b$.

To bound the fine-scale distinguishability, decompose an arbitrary state q of T^+ as $q = \sum_j w_j q_j$ with weights at most w, and let $f_j = a_j(b + (1-b)e_j) + \sum_C \mu_{j,C}\, \delta_{n(C)} 1_C$ be effects whose sum is an effect, where $n(C)$ is the level of C, $a_j + \sum_C \mu_{j,C} \le 1$, and all coefficients are nonnegative. Since $\sum_j f_j \le 1$, the coefficients satisfy $b \sum_j a_j \le 1$, so that the floored parts contribute at most w/b to $\sum_j w_j\,(f_j \cdot q_j)$. Fix L; integrating $\sum_j f_j \le 1$ against the fair coin p, under which every cell of level at most L has mass at least 2^{-L}, shows that the total coefficient $\sum_j \sum_{n(C) \le L} \mu_{j,C} \delta_{n(C)}$ is at most 2^L, so that the cells of levels at most L contribute at most $w 2^L$, while the cells of higher levels contribute at most $\sup_{n>L} \delta_n$. These bounds are uniform in q, and letting $w \to 0$ and then $L \to \infty$ gives $D(T^+) = 0$.

For the lower bound, suppose that the empty partial simulation belongs to $\mathcal{C}_2(r)$, fix $\eta > 0$, and let (H, γ, θ) in $\mathcal{C}_1(r)$ leave at most η of the mass of p outside H. Let an extension (H', γ', θ') meet a demand for a stage containing the laws p_C with $C \in \mathcal{Q}_n$, the effects $\delta_n 1_C$ and $b + (1-b)1_C$, and the constant 1, with defect at most δ', and put $q_C = \theta'(\delta_n 1_C)/\delta_n$ for $C \in \mathcal{Q}_n$. Since θ' extends linearly to the span of the effects, by the argument of [2, Lemma 2.1], the q_C are nonnegative with $\sum_C q_C = \theta'(1) \le 1$, and the response to $b + (1-b)1_C$ is $b\,\theta'(1) + (1-b)q_C$. Averaging the second inequality of Definition 4.1 over the laws p_C, at which these effects have probability 1, and arguing as in Lemma 6.1, gives

$$1 \le b + (1-b)\Big(\frac{|H|}{2^n} + \eta\Big) + r + \delta'.$$

Letting n grow and η, δ' shrink gives $r \ge 1-b$, so that $\varepsilon^{(2)}(T^+) \ge 1-b$, and Theorem 5.4 gives the rest. □

In every system computed here the classical excess is determined at the second level, $\varepsilon(A, \Delta_{\mathbb{N}}) = \varepsilon^{(2)}(A)$. This is a statement about values: later levels may still decide feasibility at the endpoint or remove other partial simulations, and the stabilization of the families $\mathcal{C}_\alpha(r)$ is a separate matter.

8 The refutation game

In [2, Section 6], simulations by $\Delta_{\mathbb{N}}$ are realizations of a type over the standard dual pair, in first-order continuous logic, where compactness decides approximate satisfiability.

Here countable additivity is written into the theory itself. For countable theories of closed conditions in real-valued infinitary logic, the infinitary extension of Keisler's model theory for real-valued structures [7], Aubone and Iovino [6] prove a refutation theorem that takes the place of compactness. In their model existence game, Olivia claims that a theory has no models and Josh attempts to describe one, committing the values of sentences to closed intervals with rational endpoints and answering Olivia's challenges; the game organizes the model existence theorem of Aubone, Dueñez and Iovino [5] and follows Väänänen's game for two-valued logic [8]. A countable theory of closed conditions has no models exactly when Olivia has a winning strategy; the plays that follow such a strategy form a well-founded tree whose leaves are finite contradictions; the tree branches infinitely only where Josh answers a challenge to a countable supremum or infimum by naming a number; and when the theory contains no countable suprema or infima, the tree is finite. In this section we write the simulations of A by $\Delta_{\mathbb{N}}$ with error at most r as the models of such a theory, and we translate the defeat of exact simulation in Theorem B(1) into a winning strategy for Olivia.

The theory. Let A be a system of countable dimension with an exhaustion by stages. Choose nested bases: an affine basis of the affine hull of Ω_1 consisting of states of A_1, extended by states of A_2 to an affine basis of the affine hull of Ω_2, and so on, and likewise nested linear bases of the spans of the sets E_n consisting of effects of the stages. Let $\Omega^{\mathbb{Q}}$ and $E^{\mathbb{Q}}$ be the countable sets of states and effects of the stages whose coordinates with respect to these bases are rational. These sets are nested along the stages, they are dense in each stage, since the rational points of the relative interior of a convex set in a finite-dimensional space are dense in it, and each of their points has a unique expansion in the common bases, so that the affine and linear relations of one stage hold in every later stage. The vocabulary L_A consists of nullary predicate symbols, each with a priori bound 1: $M_{\omega,x}$, the mass of ω at the ontic state x; $R_{e,x}$, the response of e at x; $Q_{e,\omega,x}$, their product; and $S_{\omega,N}$ and $P_{e,\omega,N}$, the partial sums over $x < N$ of the masses of ω and of the products, for $\omega \in \Omega^{\mathbb{Q}}$, $e \in E^{\mathbb{Q}}$ and $x, N \in \mathbb{N}$. Since all symbols are nullary, a structure for L_A is an assignment of numbers to them, and no witnesses or equality are involved. The theory $T_r(A)$ consists of the following closed conditions, a bound by a real number being read as the set of bounds by the rationals beyond it.

(T1) The masses $M_{\omega,x}$ and the products $Q_{e,\omega,x}$ respect the rational affine relations among the states of $\Omega^{\mathbb{Q}}$ within a stage, and the responses $R_{e,x}$ and the products $Q_{e,\omega,x}$ respect the rational linear relations among the effects of $E^{\mathbb{Q}}$ within a stage.

(T2) $M_{\omega,x} \geq 0$, $R_{e,x} \geq 0$, and $0 \leq Q_{e,\omega,x} \leq \min\{M_{\omega,x}, R_{e,x}\}$.

(T3) $Q_{e,\omega,x} = R_{e,x}M_{\omega,x}$, the product being written as a rational combination of two squares, and each square as a countable supremum of affine expressions, as in [6, Lemma 5.1].

(T4) $S_{\omega,0} = P_{e,\omega,0} = 0$, $S_{\omega,N+1} = S_{\omega,N} + M_{\omega,N}$ and $P_{e,\omega,N+1} = P_{e,\omega,N} + Q_{e,\omega,N}$.

(T5) $S_{\omega,N} \leq 1$ for every N, and $\bigvee_N S_{\omega,N} \geq 1$.

(T6) $P_{e,\omega,N} \leq e \cdot \omega + r$ for every N, and $\bigvee_N P_{e,\omega,N} \geq e \cdot \omega - r$.

The partial sums are named separately because a countable supremum requires a common a priori bound on its members [6, Section 5.1].

Proposition 8.1. *The models of $T_r(A)$ correspond to the univalent simulations of A by $\Delta_{\mathbb{N}}$ with error at most r: a model determines the simulation that gives each state ω the masses $M_{\omega,x}$ and each effect e the responses $R_{e,x}$, extended from $\Omega^{\mathbb{Q}}$ and $E^{\mathbb{Q}}$ to all states and effects, and every such simulation arises in this way.*

Proof. In a model, (T1), (T2) and (T5) make $\omega \mapsto (M_{\omega,x})_x$ a map, affine with rational coefficients, from $\Omega^{\mathbb{Q}}$ to probability vectors on $\mathbb{N}$; in each stage it is determined by its values at the finitely many generators, and extends to an affine map on Ω_n whose values are probability vectors, since nonnegativity and total mass pass from a dense set to its closure. The same holds for the responses, which by (T1), (T2) and (T3) lie in $[0,1]$ and satisfy $Q_{e,\omega,x} = R_{e,x}M_{\omega,x}$; the response to the zero effect is 0 by the rational linear relations. By (T4) and (T6), $|\sum_x R_{e,x}M_{\omega,x} - e\cdot\omega| \le r$ for $\omega \in \Omega^{\mathbb{Q}}$ and $e \in E^{\mathbb{Q}}$, and within a stage the simulated probability depends continuously on the coordinates of ω and e, because the masses of the finitely many generators have uniformly small tails; the bound therefore holds for all states and effects. Conversely, a simulation with error at most r assigns to the symbols of L_A values that satisfy every axiom. □

Olivia against an exact simulation of the Lipschitz system. The refutation below uses the axioms (T1), (T2), (T4), (T5) and (T6) only, and in particular never the products (T3) beyond the bounds of (T2). We use one consequence of the refutation theorem of Aubone and Iovino: if a position of the game, together with finitely many axioms without countable suprema, is satisfied by no structure, then Olivia has a winning strategy from that position whose tree is finite. This follows from the last clause of their theorem, applied to the finite theory formed by these axioms and the commitments of the position that contain no countable suprema, since every play from the position is a play of the game for that finite theory in which Josh has made at least the commitments it requires.

Proposition 8.2. *Let A' be the Lipschitz Cantor system, with the exhaustion by depth and nested bases chosen in the rational spans of the conditional laws p_C and of the cylinder indicators 1_C, so that these laws and the effects $2^{-n}1_C$ have rational coordinates. Olivia has a winning strategy in the game for $T_0(A')$ with the following form, along every play in which Josh has not already lost; all indices are nonnegative integers.*

(1) *Her first countable challenge concerns the total mass of the fair coin; Josh answers it by naming a number N_0.*

(2) *She then chooses a depth n that depends only on N_0, and makes 2^n further countable challenges, one for each cylinder of depth n, which Josh answers by naming numbers N_C.*

(3) *From there her strategy has a finite tree, determined by N_0 and the largest of the numbers N_C.*

Proof. Let p be the fair coin and, for a cylinder C of depth n, let p_C be its conditional law and $c = 2^{-n}$; these states, the effects $c1_C$ and $c(1-1_C)$, which are 1-Lipschitz, and the effect 1 have rational coordinates in the stage A'_n. Olivia names the axiom $\bigvee_N S_{p,N} \ge 1$ and then challenges the resulting commitment with precision 2^{-2}; Josh names N_0 and commits $S_{p,N_0} \in [\frac{3}{4}, 1]$. Olivia chooses n with $2^n > 4N_0$, and for each C she names the axiom $\bigvee_N P_{c1_C,p_C,N} \ge c$ and challenges it with a precision $2^{-k} < c/4$; Josh names N_C and commits $P_{c1_C,p_C,N_C} \in [c - 2^{-k}, 1]$. Let N_1 be the largest of the numbers N_0 and N_C; all the indices below are smaller than N_1.

The position, together with finitely many instances of (T1), (T2), (T4) and (T5) for indices below N_1 and the bound $S_{p,N_1} \le 1$ of (T5) at the cutoff, among them the linear relations $\sum_C R_{c1_C,x} = c\,R_{1,x}$ and $Q_{c\cdot 1,p_C,x} = Q_{c1_C,p_C,x} + Q_{c(1-1_C),p_C,x}$ and the nonnegativity of $Q_{c(1-1_C),p_C,x}$, is satisfied by no structure. Indeed, averaging the commitments over C gives $c - 2^{-k} \le \frac{1}{2^n}\sum_C \sum_{x<N_1} Q_{c1_C,p_C,x}$, using $Q \ge 0$ to pass from N_C to N_1. For $x < N_0$, the bound $Q_{c1_C,p_C,x} \le R_{c1_C,x}$ and the linear relation $\sum_C R_{c1_C,x} = c\,R_{1,x} \le c$ bound the inner sum by c, so that these terms contribute at most $cN_0/2^n < c/4$. For $N_0 \le x < N_1$, the linear relation $Q_{c\cdot 1,p_C,x} = Q_{c1_C,p_C,x} + Q_{c(1-1_C),p_C,x}$, the nonnegativity of the second term, and the relations $Q_{c\cdot 1,p_C,x} = c\,Q_{1,p_C,x} \le c\,M_{p_C,x}$ give $Q_{c1_C,p_C,x} \le c\,M_{p_C,x}$, and the affine relation $\frac{1}{2^n}\sum_C M_{p_C,x} = M_{p,x}$ bounds these terms by $c(S_{p,N_1} - S_{p,N_0}) \le c/4$. Hence $c - 2^{-k} < c/2$, which contradicts $2^{-k} < c/4$. The consequence of the refutation theorem recalled above gives Olivia a winning continuation with a finite tree. □

The strategy is the localization of Lemma 6.1 played in the game. Its tree branches infinitely at Josh's answer N_0 and at his answers N_C, but the two kinds of branching play different roles. The answer to the mass challenge determines which finite experiment will be used: the depth n, the conditional laws of the fair coin on the cylinders of that depth, and the effects $2^{-n}1_C$. The later answers determine only how many terms must be retained in checking that experiment. This suggests the invariant that should correspond to the derivative: not the rank of Olivia's tree, which counts every round, but the number of Josh's countable answers on which her choice of tests depends. The strategy just constructed uses one countable answer to select the finite experiment, and the empty partial simulation is removed at level 2; whether one is minimal among all strategies depends on a precise convention for the bookkeeping, which we leave open.

The raw number of countable challenges depends on the axioms, even when the class of models does not. For each depth n, the sentence $\bigvee_N \min_{C\in\mathcal{P}_n} P_{c1_C,p_C,N} \ge c$ is a consequence of $T_0(A')$, since the partial sums increase in N, so that the supremum of the minimum is the minimum of the suprema. In the theory enlarged by these sentences, an equivalent presentation of the same models, Olivia reaches a common cutoff for all the cylinders with one countable challenge, and the same finite contradiction gives a strategy with two countable challenges in all. This shows how bundling changes the count for the exhibited strategy; we do not compare the least raw counts of the two presentations.

9 When the first level suffices

We turn to two complements, which concern single models rather than the derivative. The classical excess is finitely determined exactly when the higher levels add nothing to the first. Two analytic hypotheses on a single almost optimal model force this.

Theorem 9.1. *Let A be a system, and let (Γ, Θ) be a univalent simulation of A by a classical system $\Delta_{(\Omega,\mathcal{F})}$ with error at most ϵ whose responses $\{\Theta(e) : e \in E_A\}$ form a totally bounded set for the supremum norm. Then for every $\delta > 0$ there is a positive integer k such that $\varepsilon(A, \Delta_k) \le \epsilon + \delta$.*

Proof. Let $f_1, \dots, f_m$ be responses forming a $\delta/3$-net of $\{\Theta(e) : e \in E_A\}$. Partition $[0,1]$ into finitely many intervals of length at most $\delta/3$, and let $\mathcal{P}$ be the finite measurable partition of Ω into the nonempty preimages, under the vector $(f_1, \dots, f_m)$, of the resulting product cells of $[0,1]^m$. On each cell of $\mathcal{P}$ every f_l oscillates by at most $\delta/3$, and therefore every response $\Theta(e)$ oscillates by at most δ. Choose a point x_P in each cell P, identify the

cells with $k = |\mathcal{P}|$ points, and put $\Gamma'(\omega)(P) = \Gamma(\omega)(P)$ and $\Theta'(e)(P) = \Theta(e)(x_P)$. The maps Γ' and Θ' preserve mixtures, $\Theta'(0) = 0$, and

$$\Big|\int \Theta(e)\, d\Gamma(\omega) - \sum_P \Gamma(\omega)(P)\Theta(e)(x_P)\Big| \le \sum_P \int_P |\Theta(e) - \Theta(e)(x_P)|\, d\Gamma(\omega) \le \delta$$

for every state ω and effect e, so that (Γ', Θ') is a univalent simulation of A by Δ_k with error at most $\epsilon + \delta$. □

Corollary 9.2. *If A has countably additive simulations with totally bounded responses and errors arbitrarily close to $\varepsilon_\sigma(A)$, then the classical excess of A is finitely determined, and $\varepsilon(A, \Delta_{\mathbb{N}}) = \varepsilon_\sigma(A) = \lim_k \varepsilon(A, \Delta_k)$.*

For the system A', the identity model on Cantor space has 1-Lipschitz responses, a totally bounded family by the Arzelà–Ascoli theorem, and Corollary 9.2 recovers $\varepsilon(A', \Delta_{\mathbb{N}}) = 0$.

For systems of countable dimension a weaker hypothesis on the responses suffices, at the price of a hypothesis on the preparations. Choose an affine basis $(v_i)_{i\in I}$ of the affine hull of the states of A, consisting of states and indexed by a finite or countable set I, so that every state ω is an affine combination $\sum_i a_i(\omega)v_i$ with finitely many nonzero coordinates, and choose positive weights $(w_i)_{i\in I}$ with sum 1. For a univalent simulation (Γ, Θ) of A by $\Delta_{(\Omega,\mathcal{F})}$, the measure $s = \sum_i w_i\Gamma(v_i)$ is a probability measure, the *reference measure* of the simulation, and every $\Gamma(\omega) = \sum_i a_i(\omega)\Gamma(v_i)$, by [2, Lemma 2.1], is absolutely continuous with respect to it.

Proposition 9.3. *Let A be a system of countable dimension, and let (Γ, Θ) be a univalent simulation of A by a classical system $\Delta_{(\Omega,\mathcal{F})}$ with error at most ϵ and reference measure s. Suppose that the responses $\{\Theta(e) : e \in E_A\}$ form a relatively compact subset of $L^1(s)$, and that the densities $\{d\Gamma(\omega)/ds : \omega \in \Omega_A\}$ are uniformly integrable with respect to s. Then for every $\delta > 0$ there is a positive integer k such that $\varepsilon(A, \Delta_k) \le \epsilon + \delta$.*

Proof. By uniform integrability, choose $\eta > 0$ such that $\Gamma(\omega)(B) < \delta/2$ for every state ω whenever $s(B) < \eta$, and put $t = \delta/2$. Let $g_1, \dots, g_m$ be responses forming an $\eta t/3$-net of the responses in $L^1(s)$, and choose a finite measurable partition $\mathcal{P}$ of Ω such that the conditional expectation $E_{\mathcal{P}}$ with respect to s satisfies $\|g_l - E_{\mathcal{P}} g_l\|_{L^1(s)} < \eta t/3$ for every l. Since $E_{\mathcal{P}}$ is a contraction of $L^1(s)$, every response satisfies $\|\Theta(e) - E_{\mathcal{P}}\Theta(e)\|_{L^1(s)} < \eta t$. Identify the cells of positive s-measure with k points, and put $\Gamma'(\omega)(P) = \Gamma(\omega)(P)$ and $\Theta'(e)(P) = s(P)^{-1}\int_P \Theta(e)\, ds$; the cells of measure zero carry no mass of any preparation. The maps Γ' and Θ' preserve mixtures, $\Theta'(0) = 0$, and Θ' takes values in $[0, 1]$. With $h = \Theta(e) - E_{\mathcal{P}}\Theta(e)$, which satisfies $|h| \le 1$, Markov's inequality gives $s(|h| > t) < \eta$, so that

$$\Big|\int \Theta(e)\, d\Gamma(\omega) - \Theta'(e) \cdot \Gamma'(\omega)\Big| = \Big|\int h\, d\Gamma(\omega)\Big| \le t + \Gamma(\omega)(|h| > t) < \delta$$

for every state ω and effect e. □

For the system S of Section 7, the identity model on Cantor space has responses that are totally bounded in $L^1(p)$: the indicators of the cells of level n are finitely many, with norms 2^{-n} tending to 0, and the closed convex hull of a norm-null sequence is compact. The densities $2^n 1_C$ of the conditional laws p_C, for $C \in \mathcal{Q}_n$, are not uniformly integrable, and the classical excess of S is 1 while $\varepsilon_\sigma(S) = 0$ (Proposition 7.3); the criterion fails exactly at its second hypothesis. For the system A the first hypothesis fails as well, since the coordinate indicators are effects of A at mutual $L^1(p)$ distance $\frac{1}{2}$.

10 Attainment at the countably additive optimum

A univalent simulation of A by $\Delta_{(\Omega,\mathcal{F})}$ is *atomic* if there is a countable set $Z \subseteq \Omega$, with $\{x\} \in \mathcal{F}$ for every $x \in Z$, on which every measure $\Gamma(\omega)$ is concentrated.

Proposition 10.1. *Let A be a system and $\epsilon \geq 0$. Then A has a univalent simulation by $\Delta_{\mathbb{N}}$ with error at most ϵ if and only if it has an atomic univalent simulation by a countably additive classical system with error at most ϵ. In particular, the classical excess of A equals $\varepsilon_\sigma(A)$ and is attained exactly when some optimal countably additive simulation of A is atomic.*

Proof. A simulation by $\Delta_{\mathbb{N}}$ is atomic, with $Z = \mathbb{N}$. Conversely, let (Γ, Θ) be atomic, with countable set Z, enumerated as a subset of $\mathbb{N}$. Put $\Gamma'(\omega)(x) = \Gamma(\omega)(\{x\})$ and $\Theta'(e)(x) = \Theta(e)(x)$ for $x \in Z$, and let $\Gamma'(\omega)$ vanish and $\Theta'(e)$ be 0 off Z. These maps preserve mixtures, $\Theta'(0) = 0$, and $\Theta'(e) \cdot \Gamma'(\omega) = \int \Theta(e)\, d\Gamma(\omega)$. □

The qualifier matters, since attainment of the classical excess at a value above $\varepsilon_\sigma(A)$ is a different matter: the system A attains its classical excess 1, since every simulation of A by $\Delta_{\mathbb{N}}$ has error 1, while its countably additive optimum 0 is realized by no atomic simulation. For the system A', the identity model on Cantor space is exact, its preparations are not atomic, and by Proposition 10.1 no exact simulation is atomic.

11 Questions

Question 11.1. For a system A of countable dimension, is the classical excess determined at the second level, that is, is $\varepsilon(A, \Delta_{\mathbb{N}}) = \varepsilon^{(2)}(A)$? If not, at which countable ordinal, if any, do the level excesses stabilize?

Question 11.2. Is there an intrinsic description of $\varepsilon^{(2)}(A)$, in terms of the decompositions of the states of A and the effects that distinguish their members, that accounts for effects whose scale disappears from a matching score while their positivity still constrains the responses to other effects, as for the system T^+ of Proposition 7.4, and that reduces to the bound of Proposition 7.2 for the systems of Proposition 7.3?

Question 11.3. Do the intermediate levels depend on the exhaustion when two exhaustions are not mutually cofinal by containment?

Remark 11.4. Every finite subset of the theory $T_r(A)$ of Section 8 has a model, even at error 0 and for every system A. A finite subset names finitely many indices; give every mass, response, product and named partial sum the value 0, and for each countable supremum of (T5) or (T6) that occurs, give the value 1 to a partial-sum symbol whose index exceeds every index named, which no selected axiom ties to the others. The finite subsets of this presentation therefore do not correspond to the finite stages of A, and the passage from the finite stages to the game goes through the derivative, as in Proposition 8.2.

Question 11.5. For a positive integer k, is the empty partial simulation removed at error r by level $k+1$, that is, absent from $\mathcal{C}_{k+1}(r)$, exactly when Olivia has a winning strategy in the game for $T_r(A)$ whose choice of challenges, apart from the finite bookkeeping, depends along each play on at most k of Josh's countable answers? A positive answer requires, first, a precise notion of the bookkeeping, which Proposition 8.2 isolates in one example.

A An optimal finite simulation of the system T

We construct, for the system T of Proposition 7.3, a simulation by a finite classical system with error $\frac{1}{2}(1-b)$. Put $\delta = \frac{1}{2}(1-b)$ and $h = \frac{1}{2}(1+b)$, so that $1-h = h-b = \delta$, choose a power of two $m \geq 2$ with $mb \geq h$, and let n_0 be the largest integer with $2^{n_0} \leq m$. Write $f_C = b + (1-b)1_C$ for a cell C, so that the effects of T are convex combinations of 0, the constant b and the effects f_C. Centered indicators of cells of different levels are orthogonal in $L^2(p)$, and those of one level are subject only to the relation that their sum is 0. The only affine relations among the states p_C are therefore that p is the uniform average of the laws p_C with $C \in \mathcal{Q}_n$, for each n, and the only linear relations among the effects f_C and the constant functions are that the uniform average of the effects f_C with $C \in \mathcal{Q}_n$ is the constant $b + (1-b)2^{-n}$. Assignments that respect these relations extend to mixture-preserving maps on states and effects, and since the error is a convex function of the state and of the effect, it suffices to bound it on the states p_C and the effects f_C; the constant b is sent to b, with no error.

The ontic space is the finite product $\Omega = \mathcal{Q}_1 \times \cdots \times \mathcal{Q}_{n_0} \times \{1, \ldots, m\}$, with coordinates $X_1, \ldots, X_{n_0}, Y$, and the image of p, the only state of level 0, is the uniform law on Ω; the unit effect is sent to the constant 1. For $1 \leq n \leq n_0$ and $C \in \mathcal{Q}_n$, the image of p_C is the law under which $X_n = C$ and the other coordinates are independent and uniform, and the response to f_C is 1 where $X_n = C$ and b elsewhere. For $n > n_0$, assign the cells of $\mathcal{Q}_n$ to the values of Y, $2^n/m$ cells to each value, writing $y(C)$ for the value assigned to C; the image of p_C is the law under which $Y = y(C)$ and the coordinates $X_1, \ldots, X_{n_0}$ are independent and uniform, and the response to f_C is h where $Y = y(C)$ and

$$l_n = \frac{m\big(b + (1-b)2^{-n}\big) - h}{m-1}$$

elsewhere. At every level the images of the laws p_C average to the uniform law, and the responses to the effects f_C average, at every point, to $b + (1-b)2^{-n}$: for $n \leq n_0$ this is $2^{-n}(1 + (2^n - 1)b)$, and for $n > n_0$ it is $h/m + (1 - 1/m)l_n$. The inequalities $mb \geq h$ and $2^n \geq 2$ give $0 \leq l_n \leq h$, and $l_n - b = (m(1-b)2^{-n} - \delta)/(m-1) \geq -\delta$.

It remains to compare the simulated probabilities with the true ones, which are 1 for f_C at p_C, b for f_C at $p_{C'}$ with $C' \neq C$ of the same level, and $b + (1-b)2^{-n}$ for f_C with $C \in \mathcal{Q}_n$ at p or at a law $p_{C'}$ of another level. Whenever the coordinate that carries the response to f_C is uniform under the image of the state, the simulated probability is the average $b + (1-b)2^{-n}$, which is exact; this covers the state p, the effects of levels $n \leq n_0$ at states of other levels, and the effects of levels $n > n_0$ at states of levels at most n_0. On the levels $n \leq n_0$ the simulated probabilities within a level are exact. On a level $n > n_0$, the simulated probability at p_C is h, with error δ, and at $p_{C'}$ with $C' \neq C$ it is h or l_n, with error $h - b = \delta$ or $|l_n - b| \leq \delta$. For f_C with $C \in \mathcal{Q}_n$ at a law of another level greater than n_0, the simulated probability is h or l_n, with deviations $\delta - (1-b)2^{-n}$ and $(\delta - (1-b)2^{-n})/(m-1)$ from the true value, both at most δ. The simulation therefore has error δ, and by Proposition 7.2 it is optimal.

References

[1] Lorenzo Catani, Thomas D. Galley, and Tomáš Gonda. "Resource-theoretic hierarchy of contextuality for general probabilistic theories". Quantum **10**, 2077 (2026). arXiv:2406.00717.

[2] José Iovino. "Compactness and the classical excess" (2026). Preprint.

[3] Michael Makkai. "On the model theory of denumerably long formulas with finite strings of quantifiers". Journal of Symbolic Logic **34**, 437–459 (1969).
[4] H. Jerome Keisler. "Model theory for infinitary logic". Volume 62 of Studies in Logic and the Foundations of Mathematics. North-Holland. Amsterdam (1971).
[5] Olivia Aubone, Eduardo Dueñez, and José Iovino. "Model existence and Scott's theorem for infinitary real-valued logic: a topological perspective" (2026). arXiv:2611.02247.
[6] Olivia Aubone and José Iovino. "Every obstruction to noncontextual hidden variables is finite: well-founded refutations in infinitary real-valued logic" (2026). arXiv:2611.05656.
[7] H. Jerome Keisler. "Model theory for real-valued structures". In José Iovino, editor, Beyond First-Order Model Theory, Vol. 2. Pages 43–101. CRC Press (2023). arXiv:2005.11851.
[8] Jouko Väänänen. "Models and games". Volume 132 of Cambridge Studies in Advanced Mathematics. Cambridge University Press. Cambridge (2011).