

Comment on ‘The axiom of choice and the no-signalling principle’

Martti Karvonen

Department of Computer Science, University College London
martti.karvonen@ucl.ac.uk

Abstract

The main claim of [1] is that “functional (deterministic) no-signalling resources can be stronger than probabilistic ones” in a certain nonlocal game on a Bell scenario with countably many parties. We disagree and argue that (i) under standard definitions, deterministic no-signalling resources are always probabilistic no-signalling resources; (ii) the deterministic strategy considered in [1] can be promoted to a genuinely probabilistic strategy with similar properties and (iii) a key step in the derivation in [1], claimed to hold for all no-signalling strategies, implicitly assumes measurability, leaving a gap in the argument. We propose measurability assumptions which we conjecture would make this derivation rigorous. Taken together, the phenomenon highlighted in [1] is best understood as a difference between measurable and non-measurable no-signalling resources.

1 A brief summary of [1]

In [1] the authors consider a Bell scenario with countably many parties. Modelled on a certain hat puzzle [2], a referee first chooses one of two hat colours for each party, resulting in an infinite binary string. Each party then receives as their input the suffix listing the colours of all later parties and must guess their own hat colour without any communication with the others. Intuitively, if the colours are chosen independently and uniformly at random, one gains no information from the colours of other hats and thus cannot do any better or worse than a random guess, winning half of the time. Building on this intuition, [1] claims a proof that the ratio of players that win converges to one half almost surely, and that this holds for all no-signalling strategies. However, using the axiom of choice, [1] constructs a deterministic no-signalling strategy which guarantees, for every input, that all but finitely many of the players win.

This is interpreted in [1] as showcasing a difference between deterministic and probabilistic no-signalling. In particular, [1] claims that some deterministic no-signalling strategies are not probabilistic no-signalling strategies. We disagree and instead believe that the phenomenon highlighted by [1] is best understood as a difference between measurable and non-measurable strategies.

We support this viewpoint in three ways. First, we prove that, under standard definitions, also agreed upon by the authors of [1], deterministic no-signalling resources are always probabilistic no-signalling resources (albeit of a trivial kind). In particular, the deterministic strategy constructed in [1] satisfies the probabilistic no-signalling principle as usually understood. Second, this deterministic strategy can be used to construct a genuinely probabilistic no-signalling strategy with similarly paradoxical behaviour. Third, we show that the proof in [1] that the ratio of winning players converges to one half implicitly assumes measurability when applying Azuma’s theorem [3], and we propose measurability assumptions that we conjecture would make the argument rigorous.

2 Deterministic no-signalling implies probabilistic no-signalling

In this section we demonstrate that all deterministic no-signalling resources are (degenerate) probabilistic resources, ruling out the separation claimed in [1] under standard definitions. We will explain this in a self-contained manner, assuming only a modicum of measure theory (see e.g., [4]). A Bell scenario formalizes the idea of a set of parties performing experiments on a large shared system, with each party choosing their own measurement setting for their part of it and obtaining their own outcome.

Definition 2.1. A Bell scenario $B(\mathcal{J}, \mathcal{I}, \mathcal{O})$ is specified by fixing (i) a (possibly infinite) set \mathcal{J} of parties, (ii) a (possibly infinite) set \mathcal{I} of inputs, (iii) a measurable space \mathcal{O} of outcomes. The joint measurements of $B(\mathcal{J}, \mathcal{I}, \mathcal{O})$ are given by tuples $\mathbf{x} = (x_j)_{j \in \mathcal{J}}$ in $\mathcal{I}^{\mathcal{J}}$ fixing a measurement in \mathcal{I} for each party in \mathcal{J} , and joint outcomes are elements of the measurable space $\mathcal{O}^{\mathcal{J}}$.

While our outcome set will be a finite and discrete measurable space, we still need to invoke measure theory to define joint probability distributions over joint measurements, since they live in the product $\mathcal{O}^{\mathcal{J}} = \prod_{j \in \mathcal{J}} \mathcal{O}$. We equip $\mathcal{O}^{\mathcal{J}}$ with the standard product σ -algebra generated by the cylinder sets $\prod_{j \in \mathcal{J}} X_j$ where each $X_j \subset \mathcal{O}$ and $X_j = \mathcal{O}$ for all but finitely many $j \in \mathcal{J}$. This ensures that for each $U \subset \mathcal{J}$, the projection $\pi_U: \prod_{j \in \mathcal{J}} \mathcal{O} \rightarrow \prod_{j \in U} \mathcal{O}$ onto U is measurable. When viewing elements of the product as functions, we write $\mathbf{x}|_U \in \mathcal{O}^U$ for the restriction of $\mathbf{x} \in \mathcal{O}^{\mathcal{J}}$ to U .

An empirical model¹ is an idealization of the results of such an experiment, whether as averages over multiple runs or as theoretically predicted probabilities.

Definition 2.2. An empirical model over $B(\mathcal{J}, \mathcal{I}, \mathcal{O})$ consists of a family $e = (e_{\mathbf{x}})_{\mathbf{x} \in \mathcal{I}^{\mathcal{J}}}$ where for each joint input $\mathbf{x} \in \mathcal{I}^{\mathcal{J}}$ we have a probability measure $e_{\mathbf{x}}$ on the space $\mathcal{O}^{\mathcal{J}}$ of joint outcomes. We write $e: B(\mathcal{J}, \mathcal{I}, \mathcal{O})$ to denote that e is an empirical model over $B(\mathcal{J}, \mathcal{I}, \mathcal{O})$.

Arbitrary empirical models are too general in certain physically motivated situations. If the parties are space-like separated, then choices made by some parties cannot causally influence the outcomes of others, assuming relativity. It is then natural to postulate the *no-signalling conditions* which assert that the distribution for a given set of parties does not depend on the inputs of the parties in the complement. Formally, we note that each restriction $|_U: \mathcal{O}^{\mathcal{J}} \rightarrow \mathcal{O}^U$ is measurable. As a result, for any probability measure d on the space $\mathcal{O}^{\mathcal{J}}$ of joint outcomes, there is a well-defined push-forward probability measure on \mathcal{O}^U along the restriction to U which is usually known as the marginalization of d to U . We will overload notation and denote the resulting probability measure by $d|_U$. With this out of the way, we can formally define no-signalling.

Definition 2.3. An empirical model $e: B(\mathcal{J}, \mathcal{I}, \mathcal{O})$ is a no-signalling empirical model, if for all joint inputs $\mathbf{x}, \mathbf{y} \in \mathcal{I}^{\mathcal{J}}$ we have

$$e_{\mathbf{x}}|_{[\mathbf{x}=\mathbf{y}]} = e_{\mathbf{y}}|_{[\mathbf{x}=\mathbf{y}]} \quad (2.4)$$

where $[\mathbf{x} = \mathbf{y}]$ denotes the set $\{j \in \mathcal{J} : \mathbf{x}_j = \mathbf{y}_j\}$ of indices where \mathbf{x} and \mathbf{y} agree.

We emphasize that these definitions of (no-signalling) empirical models are standard: for instance, they agree with those given in [5] which directly generalize the well-established Abramsky–Brandenburger framework for contextuality [6] into the infinitary setting. Moreover, in an email exchange about these issues when [1] first appeared on arXiv, the corresponding author confirmed that “this is the natural definition of (probabilistic) no-signaling and it is indeed the one that we adopted throughout.” [7].

¹Also known as a correlation, behaviour, or a strategy, depending on the source.

We stress that, due to the measurability of the restrictions, the probability measures appearing in (2.4) are always well defined, no matter how pathological e is otherwise. This is not to say that every empirical model is always no-signalling, only that the question of whether it is no-signalling is always well posed. We now define the class of local empirical models.

Definition 2.5. Let μ be a probability measure on $\mathcal{O}^{\mathcal{J} \times \mathcal{I}}$. Define an empirical model $e_\mu = (e_{\mu, \mathbf{x}})_{\mathbf{x} \in \mathcal{I}^{\mathcal{J}}}$ by setting $(e_{\mu, \mathbf{x}}) := \mu|_{\mathbf{x}}$, where we write $\mu|_{\mathbf{x}}$ for the push-forward of μ along the restriction to the graph $\{(j, \mathbf{x}_j) : j \in \mathcal{J}\} \subset \mathcal{J} \times \mathcal{I}$ of $\mathbf{x} \in \mathcal{I}^{\mathcal{J}}$. We say that an empirical model e is local if $e = e_\mu$ for some probability measure μ on $\mathcal{O}^{\mathcal{J} \times \mathcal{I}}$.

Theorem 2.6. All local empirical models are no-signalling.

Proof. Follows directly from functoriality of taking push-forward measures, i.e., the fact that pushing a measure forward twice is the same as pushing it forward once along the composite measurable mapping. Explicitly, let e_μ be local, $\mathbf{x}, \mathbf{y} \in \mathcal{I}^{\mathcal{J}}$ be arbitrary, and compute as follows:

$$e_{\mu, \mathbf{x}}|_{[\mathbf{x}=\mathbf{y}]} = (\mu|_{\mathbf{x}})|_{[\mathbf{x}=\mathbf{y}]} = \mu|_{[\mathbf{x}=\mathbf{y}]} = (\mu|_{\mathbf{y}})|_{[\mathbf{x}=\mathbf{y}]} = e_{\mu, \mathbf{y}}|_{[\mathbf{x}=\mathbf{y}]}. \quad \square$$

In particular, if μ is given by the Dirac measure $\delta(f)$ on some function $f: \mathcal{J} \times \mathcal{I} \rightarrow \mathcal{O}$, then it always induces a no-signalling empirical model. Similarly, finite and countable convex combinations of such models are no-signalling, i.e., if μ_i for $i \in \mathbb{N}$ are Dirac measures and r_i for $i \in \mathbb{N}$ are nonnegative reals summing to 1, then $\mu := \sum_{i \in \mathbb{N}} r_i \mu_i$ induces a no-signalling empirical model.

We now turn our attention to the *functional no-signalling principle* discussed in [1, Equation 1.3]. In this setting, one considers a function $f: \mathcal{I}^{\mathcal{J}} \rightarrow \mathcal{O}^{\mathcal{J}}$ from joint inputs to joint outputs, written as $f(\mathbf{x}) = (f_j(\mathbf{x}))_{j \in \mathcal{J}}$ for $\mathbf{x} = (x_j)_{j \in \mathcal{J}}$ in $\mathcal{I}^{\mathcal{J}}$. Any such function f induces an empirical model $e(f): B(\mathcal{J}, \mathcal{I}, \mathcal{O})$ by setting the distribution $e(f)|_{\mathbf{x}}$ on $\mathcal{O}^{\mathcal{J}}$ to be the Dirac measure on $f(\mathbf{x})$.

Such a function is said to satisfy the *functional no-signalling principle* if the j th output depends only on the j th input, which can be expressed by saying that for all $j \in \mathcal{J}$ and all $\mathbf{x}, \mathbf{y} \in \mathcal{I}^{\mathcal{J}}$, if $x_j = y_j$, then $f_j(\mathbf{x}) = f_j(\mathbf{y})$. In that case the function $f: \mathcal{I}^{\mathcal{J}} \rightarrow \mathcal{O}^{\mathcal{J}}$ in fact determines a function $\hat{f}: \mathcal{J} \times \mathcal{I} \rightarrow \mathcal{O}$, defined by $\hat{f}(j, i) = f_j(\mathbf{x})$, where \mathbf{x} is an arbitrary element of $\mathcal{I}^{\mathcal{J}}$ satisfying $x_j = i$. As f satisfies the functional no-signalling principle, \hat{f} is indeed a well-defined function, as the choice of \mathbf{x} does not affect the end result. Moreover, $e(f) = e_{\delta(\hat{f})}$ i.e., $e(f)$ is induced by the Dirac measure on \hat{f} , so that $e(f)$ is local. Therefore, Theorem 2.6 immediately gives the following result.

Theorem 2.7. If $f: \mathcal{I}^{\mathcal{J}} \rightarrow \mathcal{O}^{\mathcal{J}}$ satisfies the functional no-signalling principle, it induces a no-signalling empirical model $e(f)$ on $B(\mathcal{J}, \mathcal{I}, \mathcal{O})$.

Thus all functional no-signalling resources are no-signalling probabilistic resources (albeit of a degenerate kind), ruling out the claimed separation under standard definitions.

3 A probabilistic counterexample

We now explain the setting of [1] in more detail. The main construction is based on an ‘infinite hat puzzle’ [2]: there is a player for each natural number, and a referee equips each player with a hat with one of two colours $\{0, 1\}$. The j th player sees the colours of the hats of all later players (i.e., the i th hat for $i > j$) but no other hats. The players cannot communicate after the setup but may agree on a strategy beforehand, possibly using shared randomness, and each must guess the colour of their own hat based on what they see and the agreed strategy, with the referee recording for each player whether they won (by guessing correctly) or lost.

We model the game on the Bell scenario $B(\mathbb{N}, \{0, 1\}^{\mathbb{N}}, \{0, 1\})$ with parties given by \mathbb{N} , inputs² by $\{0, 1\}^{\mathbb{N}}$ i.e., sequences of colours, and two outputs $\{0, 1\}$. The game itself can be modelled as follows: the referee first picks each hat independently and uniformly at random, i.e., samples a string $\mathbf{x} = (x_0, x_1, x_2, \dots)$ from the uniform product distribution on $\{0, 1\}^{\mathbb{N}}$. The interpretation is that $x_j \in \{0, 1\}$ is the colour of the j th hat. The referee then sends each player an input: the j th player is given the input $(x_{j+1}, x_{j+2}, \dots)$, i.e., the colours of the hats of all the later players. They then output their guess a_j , and the result r_j of the j th player is then given by

$$r_j = \begin{cases} 1 & \text{if } a_j = x_j \\ 0 & \text{otherwise.} \end{cases}$$

Assuming the players have no limits on their memory or computational abilities, they can counter-intuitively construct a strategy that guarantees that all but finitely many of them win, using the *axiom of choice*. Consider the equivalence relation \sim on $\{0, 1\}^{\mathbb{N}}$ of having equal tails after some threshold: formally, $\mathbf{x} \sim \mathbf{y}$ if there is a $n \in \mathbb{N}$ such that $x_j = y_j$ for all $j > n$. Using the axiom of choice, the players fix one representative in each equivalence class and encode these choices in a function $g: \{0, 1\}^{\mathbb{N}} \rightarrow \{0, 1\}^{\mathbb{N}}$ sending each \mathbf{x} to the chosen representative of its class. The strategy of the j th player is then as follows: given the input $(x_{j+1}, x_{j+2}, \dots)$, they prepend it with $j + 1$ zeroes, obtaining a string $\mathbf{y}_j := (0, \dots, 0, x_{j+1}, x_{j+2}, \dots) \in \{0, 1\}^{\mathbb{N}}$, and set their guess a_j to be $g(\mathbf{y}_j)_j$. By construction, for each $j \in \mathbb{N}$ we have $\mathbf{y}_j \sim \mathbf{x}$ and therefore $g(\mathbf{y}_j) = g(\mathbf{x})$. As $g(\mathbf{x}) \sim \mathbf{x}$, this means that after some threshold $n(\mathbf{x}) \in \mathbb{N}$, all players guess correctly.

This strategy has each player deterministically choose their outcome based only on their own input, thus inducing a function $\mathbb{N} \times \{0, 1\}^{\mathbb{N}} \rightarrow \{0, 1\}$ and hence by Theorem 2.6 a no-signalling empirical model on $B(\mathbb{N}, \{0, 1\}^{\mathbb{N}}, \{0, 1\})$. Now, [1] summarizes their main result as follows: “Hence, imposing probabilistic NS, the ratio of players that win the game over the total converges to 1/2 almost surely.” However, the above strategy gives a counterexample. This counterexample is in fact discussed in [1] but interpreted differently: it is treated as a functional no-signalling strategy which fails to be a probabilistic NS model, since in the presence of non-measurable sets (such as the set of representatives of the equivalence classes above) some probabilities may fail to exist. While it is indeed true that in the presence of non-measurable sets not all conceivable sets are assigned a well-defined probability, the presence of some non-measurable sets does not imply that all relevant probabilities fail to exist. Indeed, as discussed earlier, the restrictions $\{0, 1\}^{\mathbb{N}} \rightarrow \{0, 1\}^J$ being measurable guarantees that all probabilities required to express the (probabilistic) no-signalling equations exist as soon as one has an empirical model, no matter how many non-measurable sets were used to construct it. Hence this strategy does define a probabilistic no-signalling empirical model, providing a direct counterexample to their claim.

To emphasize that this is not restricted to deterministic strategies, let g_j be obtained from g by swapping the first j outputs and keeping the rest as in g . Let e_j be the empirical model induced by g_j , and define $d := \frac{1}{2}e_0 + \frac{1}{2}e_1$ and $e := \sum_{j \in \mathbb{N}} \frac{1}{2^j} e_j$. These models are no-signalling by Theorem 2.6: the behaviour of d is uniformly random at the first party (and deterministic otherwise), while the behaviour of e is genuinely probabilistic for every party (albeit it tends to deterministic behaviour exponentially fast). Moreover, for each fixed joint input, the strategy d guarantees that all but

²[1] works with inputs given by the unit interval $I = [0, 1]$. Taking binary expansions maps the unit interval to $\{0, 1\}^{\mathbb{N}}$, and yields an isomorphism of probability spaces, up to null sets corresponding to numbers with two binary expansions and the corresponding binary sequences, when the unit interval is endowed with the uniform distribution (Lebesgue measure) and $\{0, 1\}^{\mathbb{N}}$ with the (completion of the) product of infinitely many coin flips. Therefore, we can view this choice as just a matter of notational convenience. Moreover, working with the slightly less familiar measurable space $\{0, 1\}^{\mathbb{N}}$ discourages computing integrals without first verifying that they are well defined.

finitely many of the players win, and for e this happens almost surely (by an application of the first Borel–Cantelli lemma). It follows that neither satisfies the bounds derived in [1]. These two genuinely probabilistic models show that the findings of [1] do not demonstrate a separation between deterministic and probabilistic no-signalling strategies under standard definitions.

4 Missing measurability assumption

In [1] it is claimed that, under probabilistic no-signalling, the ratio of winning players converges to one half almost surely. We have now exhibited a probabilistic counterexample, so this claim does not hold for all no-signalling probabilistic empirical models. We next identify a gap in their derivation, and propose a stronger assumption which we conjecture to lead to a corrected proof, applying only to a subclass of no-signalling empirical models.

The key gap we have identified is the following: at a crucial juncture the argument applies Azuma’s theorem [3] which concerns certain sums of random variables. Here, the candidate random variables S_n are given by sums $S_n = \sum_{j=1}^n s_j$ of “success random variables” s_j where $s_j = 1$ if the j th player wins and $s_j = -1$ otherwise.

Recall that, formally, a real-valued random variable is a measurable function into the reals. The domain of this function s_j is never explicitly defined in [1], but it is reasonable to take it to be the space $\{0, 1\}^{\mathbb{N}}$, since that is the sample space for the referee’s random assignments of hats (or inputs).

We now argue that these “random variables” need not be measurable functions, in which case the assumptions of Azuma’s theorem are not satisfied. The functions S_n are measurable iff the functions s_j are, so we focus on the latter. We first consider the situation when players are using the deterministic strategy constructed above using the function $g: \{0, 1\}^{\mathbb{N}} \rightarrow \{0, 1\}^{\mathbb{N}}$ obtained via the axiom of choice. Then s_j is given by

$$s_j(\mathbf{x}) := \begin{cases} 1 & \text{if } g(\mathbf{x})_j = x_j \\ -1 & \text{otherwise.} \end{cases}$$

As this function takes only two values, we see that s_j is measurable iff the set $G_j := \{\mathbf{x} \in \{0, 1\}^{\mathbb{N}} : g(\mathbf{x})_j = x_j\}$ of inputs for which the j th party wins when following the strategy induced by g is measurable. If the sets G_j were measurable, so would be the set $\bigcap_{j=1}^{\infty} G_j = \{\mathbf{x} \in \{0, 1\}^{\mathbb{N}} : g(\mathbf{x}) = \mathbf{x}\}$ of inputs where all players win. This coincides with the set of chosen representatives, and as a direct analogue of the well-known Vitali set on the reals can be shown to be non-measurable similarly.

Although the point is clearest for deterministic strategies, the mixed model $d := \frac{1}{2}e_0 + \frac{1}{2}e_1$ above exhibits the same issues for all $j > 1$. For $e := \sum_{j \in \mathbb{N}} \frac{1}{2^j} e_j$, one can also show that the (now probabilistic) maps $s_j: \{0, 1\}^{\mathbb{N}} \rightarrow \{-1, 1\}$ are also not measurable, although this takes some more work. Interestingly, none of these examples are counterexamples to the derived bound in the sense of violating it: rather, they fall outside its scope, as the probabilities the bound concerns do not exist for these strategies.

5 Conclusion

In conclusion, the main phenomenon highlighted in [1] is not about deterministic versus probabilistic strategies, nor about the no-signalling principle itself. Rather, the issue is that an empirical model may satisfy the no-signalling principle yet be pathological in the sense of being non-measurable.

To formalize this, recall that an empirical model $e : B(\mathcal{J}, \mathcal{I}, \mathcal{O})$ is defined as a family $e = (e_{\mathbf{x}})_{\mathbf{x} \in \mathcal{I}^{\mathcal{J}}}$. Equivalently, it can be thought of as a function sending each joint input $\mathbf{x} \in \mathcal{I}^{\mathcal{J}}$ to the distribution $e_{\mathbf{x}}$ over joint outcomes. It is known that the set $\mathcal{D}(X)$ of distributions on a measurable space X can itself be made into a measurable space, so that measurable maps $X \rightarrow \mathcal{D}(Y)$ correspond to stochastic maps (Markov kernels) $X \rightarrow Y$ [8]. We propose to call an empirical model $e : B(\mathcal{J}, \mathcal{I}, \mathcal{O})$ *measurable* if it is measurable when viewed as a function from joint inputs to distributions over outcomes with respect to the product σ -algebras fixed above. We conjecture that the results of [1] apply to all measurable empirical models, and hence they contribute another example to the long list of differences between measurable and non-measurable behaviours. Such measurability assumptions seem perfectly reasonable in physical models, rule out the pathological examples discussed above, and they hold trivially when the sets of parties and inputs are finite. However, it is crucial in science to state one’s assumptions explicitly.

Acknowledgements. Martti Karvonen acknowledges useful conversations with Rui Soares Barbosa and support from the EPSRC grant EP/V040944/1 Resources in Computation.

References

- [1] Baumeler A, Dakić B, Del Santo F. 2025 The axiom of choice and the no-signalling principle. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **481**. (10.1098/rspa.2024.0601)
- [2] Hardin CS, Taylor AD. 2013 *The Mathematics of Coordinated Inference: A Study of Generalized Hat Problems*. Springer International Publishing. (10.1007/978-3-319-01333-6)
- [3] Azuma K. 1967 Weighted sums of certain dependent random variables. *Tohoku Mathematical Journal* **19**. (10.2748/tmj/1178243286)
- [4] Billingsley P. 1979 *Probability and Measure*. Wiley Series in Probability and Mathematical Statistics. New York: Wiley.
- [5] Barbosa RS, Douce T, Emeriau PE, Kashefi E, Mansfield S. 2022 Continuous-Variable Nonlocality and Contextuality. *Communications in Mathematical Physics* **391**, 1047–1089. (10.1007/s00220-021-04285-7)
- [6] Abramsky S, Brandenburger A. 2011 The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics* **13**, 113036.
- [7] del Santo F. 2022 Re: AC and NS. Private communication (email). Email to Martti Karvonen, 25 July 2022.
- [8] Giry M. 1982 p. 68–85. In *A categorical approach to probability theory*, p. 68–85. Springer Berlin Heidelberg. (10.1007/bfb0092872)