8

Quantum Correlations

In this chapter, we use BST to analyze the phenomenon of quantum mechanical (QM) correlations. Our hope is that our BST analysis, based on rigorously developed notions of transitions, propensities, and funny business, sheds some light on these puzzling phenomena. We conceive of this project as a modest one: we do not aim to produce a general "BST interpretation of QM" or anything similar. We study how the machinery of BST can be applied to the modeling of the strange correlations that QM predicts and which have been so well confirmed experimentally, and what we can learn from attempts at getting rid of the correlations via extended models.

Quantum mechanical correlations seem to be ideally suited for description and analysis in the BST framework. They concern events occurring in space and time, such as the selection of measurement parameters, measurements, or various possible detection events. The spatio-temporal features of the set-up matter, as some of the events involved are space-like related. And the correlations involve modal issues of possibility and impossibility, or of grades of possibility. We have already mentioned purely modal QM correlations when motivating the investigation of what we have called modal funny business (MFB) in BST (Chapter 5.1) via the quantum-mechanical EPR set-up. That set-up involves an entangled particle pair, leading to perfect anticorrelations between space-like separated measurements (see Figure 5.1, p. 107): some joint outcomes are impossible even though the individual outcomes are separately possible. A more complex, but also more interesting case of modal correlations is the so-called Greenberger-Horne-Zeilinger (GHZ) set-up, which involves an entangled three-particle state for which, again, some individually possible measurement outcomes are jointly impossible. Of course, this has more than a whiff of MFB about it. The GHZ set-up will play a prominent role in our discussions in Section 8.3, especially in Section 8.3.3.

Apart from *modal* correlations, and more in the focus of discussions in the philosophy of physics, QM also predicts space-like *probabilistic*

Branching Space-Times: Theory and Applications. Nuel Belnap, Thomas Müller, and Tomasz Placek,
Oxford University Press. © Oxford University Press 2022. DOI: 10.1093/oso/9780190884314.003.0008

correlations, which we have called probabilistic funny business (PFB; see Chapter 7.2.6). In such correlations, joint results of individually possible, space-like related outcomes are possible all together, but the degree to which they are possible—their joint causal probabilities or propensities—cannot be derived from the probabilities of the individual outcomes taken separately. An important set-up in this regard is the so-called Bell-Aspect set-up, which will be a focus for our discussions of PFB in Section 8.4, especially in Section 8.4.4.

Since the famous EPR paper by Einstein et al. (1935), but especially following the ground-breaking works of John Bell in the 1960s (see Bell, 1987a), quantum correlations have been analyzed in terms of the possibility of introducing an additional explanatory structure, that of the so-called hidden variables. The question from the title of the EPR paper, "Can quantum-mechanical description of physical reality be considered complete?", can be answered by studying which extensions, or completions, of the quantum-mechanical surface description of a QM correlation experiment are empirically viable. Over the years, starting with Bell's results, there have been a large number of "no go" theorems concerning the introduction of certain classes of hidden variables for various set-ups. In the BST framework, we can study the introduction of hidden variables as structure extensions that lead from a BST surface structure to an enlarged "hidden" structure that has more desirable features than the surface structure. Ideally, while the surface structure harbors MFB (or PFB), the extended structure represents the same surface phenomena, but is free from MFB (or PFB). The "no go" theorems say that this is not always possible.

In BST, we can reproduce these results in a framework representing both the spatio-temporal and the modal aspects. The BST versions of the "no go" theorems we derive show why it is impossible to explain certain cases of modal or probabilistic funny business by invoking hidden variables, and which modal, causal, propensity-related, and spatio-temporal features of our world these theorems rely on. Another significant contribution of a BST analysis, in our view, lies in the fact that we can spell out in formal detail what it means to analyze a set-up *as an experiment*, and which role this plays in the derivation of the "no go" results. It turns out that it is crucial to keep two types of indeterminism separate: indeterminism due to an experimenter's selection of measurement parameters, and indeterminism due to Nature's choice of a measurement result. BST makes room for a transparent representation of this distinction.

Our analysis in this chapter is framed in terms of QM, but our analysis of structure extensions vis-à-vis modal or probabilistic funny business is of a general nature. It can be applied to any case in which one wonders whether or not funny business is, as it were, empirically inevitable. Quantum mechanics is just the theory that gives us the strongest reasons for thinking that this may be so.

8.1 Introducing quantum correlation experiments

Put in abstract terms, a quantum correlation experiment has the following general form. The experimental set-up contains a source of systems (think of entangled pairs or n-tuples of particles) that are channeled into two or more stations, assigned to experimenters Alice (A), Bob (B), Carol (C), etc. The measurements conducted at each of the stations can be varied. In the fairly simple Bell-Aspect set-up, each experimenter is in control of two alternative settings of her measurement apparatus, 1 and 2 for Alice and 3 and 4 for Bob, and each of the possible measurements has two possible individual outcomes, '$+$' and '$-$'. For a given two-particle state, quantum mechanics predicts which joint outcomes are possible for given settings, and what the probabilities for these joint outcomes are. A single run of the experiment can thus be described by listing the selected settings together with the results obtained with these settings. So, in a set-up with just two experimenters, $a_1^+ b_3^-$ describes a run in which Alice selected setting 1, Bob selected setting 3, Alice obtained result '$+$', and Bob obtained result '$-$'. The Bell-Aspect set-up is pictured schematically in Figure 8.1.

What we have just described schematically is an experiment rather than a collection of natural happenings. The idea behind doing experiments is to pose questions to Nature by exerting control over the conditions and then observing what happens. Intervention is crucial to an experiment, yet conceptually it cannot be determined or dictated by Nature. There must be a certain independence between what happens anyway and which questions are asked in an experiment. In basic BST terms, we can incorporate the distinction between natural happenings and experimental interventions by distinguishing between two types of local choices: Some of these choices are assumed to be under experimental control, while the rest of the indeterminism in the model is due to Nature alone. This way of describing experiments

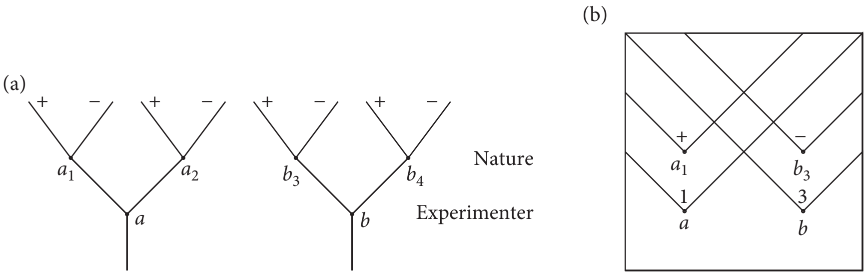

Figure 8.1 Schematic illustration of the Bell-Aspect quantum correlation experiment. Experimenters Alice (a) and Bob (b) choose settings (1 or 2 for Alice, 3 or 4 for Bob) for their experiments, which have outcomes '$+$' or '$-$'. Left: Schematic illustration of the modal splitting involved. Right: Space-time diagram for the history including the joint outcome $a_1^+ b_3^-$.

is in fact deeply ingrained in the whole quantum correlation literature. Here are some relevant quotes. Einstein et al. (1935) speak of the consequences of performing different measurements:

> We see therefore that, as a consequence of two different measurements performed upon the first system, the second system may be left in states with two different wave functions. On the other hand, since at the time of measurement the two systems no longer interact, no real change can take place in the second system in consequence of anything that may be done to the first system. (Einstein et al., 1935, p. 779)

Bell (1987b) speaks of the freedom of experimenters:

> It has been assumed that the settings of instruments are in some sense free variables—say at the whim of experimenters—or in any case not determined in the overlap of the backward light cones. Indeed without such freedom I would not know how to formulate *any* idea of local causality, even the modest human one. (Bell 1987b, p. 61)

Similar remarks can be found in papers in experimental physics. For example, Aspect et al. (1982b) mention "other choices of orientations". In our schematic set-up, we thus contrast the selection of settings, which is under experimental control, and the chancy occurrence of the measurement outcomes, which is due to Nature. Experimental control does not have to mean that the selection of settings is due to a human agent.[1] A distinctive feature of experimental control, which enters as an assumption in the derivation of results about the possibility or impossibility of certain types of hidden variables, such as the derivation of Bell-type theorems (see Section 8.4), is that the selection of settings must be independent of the physical state of the system on which the experiment is performed—including a hypothetical full state that adds hidden variables. Typically, this independence is understood as the existence of truly indeterministic processes of the selection of settings whose outcomes are, in particular, not influenced by the values of any hidden variables.[2] In contrast to the settings, the measurement outcomes may of course depend upon, or even to be dictated by, the full physical state including the values of hidden variables.

8.2 On the BST analysis of quantum correlations

Since Einstein, Podolski and Rosen's (1935) seminal paper, quantum correlations have been viewed as mysterious; they seem spooky, or "spukhaft", as Einstein had it.[3] After all, such correlations coordinate remote space-like separated measurement results, like a_1^+ and b_3^- in our schematic illustration, without the possibility of physical interaction. In some experiments, the spatial distance between the results is enormous.[4] BST has two resources to represent correlations between space-like related measurement results, modal funny business (MFB) and probabilities funny business (PFB). The first approach applies to quantum-mechanical correlations only if some joint outcomes are assigned probability zero. On that modal (MFB) approach, probability zero is interpreted as meaning that a joint outcome is impossible, while a non-zero probability signals that a joint outcome is possible.

[1] The experiment of Aspect et al. (1982a), with semi-random selections of settings produced by an optical process, was viewed as incorporating a type of experimental control that constituted a decisive improvement over earlier experiments in which the settings were fixed for whole series of runs of the experiment. Even so, direct human control seems to be the conceptual gold standard for such experiments, despite the fact that the quality of the randomness produced by humans, as assessed via statistical tests, is much inferior to other sources of randomness. For a recent large-scale experiment using direct human control for the selection of parameters in quantum correlation experiments, see Abellán et al. (2018).

[2] For a dissenting view, see Esfeld (2015).

[3] See footnote 2 in Chapter 5.

[4] As of Spring 2019, the longest distance is 144 km on the ground (Scheidl et al., 2010) and 1200 km by satellite communication (Yin et al., 2017).

Measurement events are represented as transitions, and a set-up with modal correlations can be analyzed using the resources developed in Chapter 5. The second approach applies to all cases of quantum-mechanical correlations, including, but not limited to those in which probabilities zero and one occur. On that probabilistic (PFB) approach, probabilities of joint outcomes are interpreted as (single case) propensities, which can be analyzed using the framework developed in Chapter 7.2.6.

According to quantum mechanics, remote correlations can indeed be perfect (at least in theory), meaning that the probability of the joint occurrence of two specific, individually possible results can be zero. This implies that knowing her own result, an experimenter can immediately predict a remote result. This looks as if quantum correlations could beat the limit of the speed of light that is imposed on physical interactions by the special theory of relativity. This leads to the worry of whether quantum correlations indicate a conflict between the theories of relativity and of quantum mechanics, and in particular, whether they permit superluminal signaling. The consensus is that while entanglement can be used to greatly enhance the security of communication (Colbeck and Renner, 2012), it does not provide a resource for superluminal communication. There are space-like correlations involved, but these cannot be exploited by the experimenters to send signals. More generally, the consensus appears to be that there is no logical conflict between relativity and quantum mechanics, and that the two theories live in a kind of peaceful coexistence, as argued first by Shimony (1978). Needless to say, as this very phrase is reminiscent of Cold War talk,[5] the implication is that there are decisive conflicts between the two theories, but these are not lethal, as they do not amount to a head-on contradiction. Thus, the phrase also suggests some hope of resolving the conflict via a future unified theory encompassing large- and small-scale phenomena together.[6]

[5] Shimony attributed the phrase to Chairman Khrushchev; "mirnoye sosushchestvovaniye", as it is called in Russian, was part of Soviet propaganda in the Cold War and afterward.

A key element in Shimony's (1978) peaceful coexistence strategy is to uphold one premise of Bell's theorem (the independence of measurement outcomes from the choice of remote settings), while rejecting a mathematically similar premise that remote measurement outcomes are independent from each other (for a precise statement of these premises, see Def. 8.24 for Parameter Independence and Def. 8.22 for Outcome Independence, both discussed in Section 8.4). This suffices to block the derivation of Bell's theorem, while prohibiting faster than light signalling. Given the mentioned similarity of the two premises, the move is nevertheless controversial.

[6] Given more recent developments, like the Colbeck-Renner (2011) proof of a Bell-type theorem, or purported progress with Bohmian analyses of EPR correlations (for a survey, see, e.g., Maudlin 2019, ch. 4), the idea of peaceful coexistence now looks more complicated than ever. For a recent reassessment, see Butterfield (2018).

Correlations between space-like separated events raise the question of explanation and especially of constraints on explanations that relate to "locality." If correlated events are space-like separated, a correlation cannot be explained via direct physical interactions—at least this is the consensus view. An explanation would thus have to come from an extended description of the states of the system involved, as already argued by Einstein et al. (1935). The big question is, of course, which extensions are admissible conceptually and empirically, and which kind of explanation they can provide. The initial thrust of the quest for hidden variables is often interpreted as an attempt to eliminate the indeterminism of quantum mechanics, or to reduce quantum-mechanical probabilities to epistemic phenomena.[7] Over the years, however, it has become clear that the real challenge is to explain away any supposed "non-locality" of quantum mechanics, not the theory's indeterminism. Thus, one constraint on the hidden-variable completions of quantum mechanics is that the resulting models be "local", where the informal notion of locality must be made precise in some way. A natural way to explain QM correlations locally is to postulate common causes lying in the past of correlated events, or to postulate instructions that the correlated particles carry along from the source to the measurement. Starting with the work of John Bell in the 1960s (see Bell, 1987a), it has become clear that postulating such hidden variables is not a merely metaphysical maneuver, but can have implications that can be tested empirically. Given the well-established empirical fact that observed correlations vindicate the predictions of QM,[8] the issue has become less one of extending the quantum-mechanical formalism, but of coming up with any viable account of the observed correlations at all. For such an account to count as explanatory, it has to fulfill certain structural constraints that are either physically or philosophically motivated. Crucially, they include the

[7] Einstein is often attributed with the slogan "God does not play dice", but the interpretation of this slogan is controversial. In a letter to Max Born dated December 4, 1926, Einstein writes: "Die [Quantenmechanik] liefert viel, aber dem Geheimnis des Alten bringt sie uns kaum näher. Jedenfalls bin ich überzeugt, daß *der* nicht würfelt." See Einstein et al. (1971, p.90): "[Quantum mechanics] says a lot, but does not really bring us any closer to the secret of the 'old one'. I, at any rate, am convinced that *He* is not playing at dice."

[8] A whole literature is devoted to the important question of whether the possible "loopholes" often present in actual experiments can be and have been closed. For example, if detection efficiency is too low, it is impossible to check whether the runs registered constitute a fair sample of all really possible runs. For a discussion of loopholes, see, e.g., Myrvold et al. (2019). Over the years, substantial progress in closing these loopholes has been made. Some recent experiments have been claimed to be "loophole-free" (Giustina et al., 2013; Hensen et al., 2015; Shalm et al., 2015). It appears that at least the scientific community is convinced that in these experiments, the issues of loopholes have been dealt with satisfactorily, so that the quantum-mechanical predictions are really empirically vindicated.

independence of experimenters' choices of measurement parameters from the particles' state.

The basic *local* features of BST can help to analyze the quest for local explanations. In BST, all possible events are well-defined spatio-temporal objects. If such an event is an indeterministic choice point, it has possible outcomes that occur *after* it in the BST pre-causal ordering. This basic locality is supplemented by two restricted and well-defined varieties of non-locality: BST permits modal as well as probabilistic funny business (see Chapter 5 and Chapter 7.2.6, respectively). With these resources, BST is able to resolve issues that are outside the scope of other rigorous, but purely probabilistic frameworks available in the literature.[9] The BST resources will be especially helpful when tackling the independence of experimental interventions and natural indeterminism. This independence has at least two aspects. First, an experimenter, who has a range of possible measurement parameters to choose from, should not be restricted in her choice. This might perhaps sound overly broad: undoubtedly, Nature limits our ways, and it may not be possible to construct an experimental set-up in which one can ask Nature different questions by selecting alternative parameters. The independence that is at stake here, however, is more specific: given that an apparatus is in place that affords the choice of different parameters, there must be no surreptitious or hidden limitation of the experimenters' choices in any given run of the experiment. The other aspect goes in the opposite direction: the experimenter must not be able to restrict or influence Nature's choices at remote (space-like separated) events.

The notion of independence still needs to be made formally precise in a way that accords with the general framework that is used, but here we can already formulate a template for the independence condition, which we will later call C/E independence:

Definition 8.1 (Target notion of C/E independence)**.** Let $\mathscr{W} = \langle W, < \rangle$ be a BST_{92} structure with two disjoint sets of choice points, $C, E \subseteq W$, where C represents the choices of experimenters and E represents indeterminism due to Nature. The structure is said to be *C/E independent* iff for any consistent subsets $C_0 \subseteq C$ and $E_0 \subseteq E$, the outcomes of C_0 and the outcomes of E_0 are independent in the relevant (modal or probabilistic) sense.

[9] For such accounts, see, e.g., Hofer-Szabó et al. (1999) and Pitowsky (1989).

We already said that our analysis focuses on the representation of the possible single runs of quantum correlation experiments. Each single run is a complex spatio-temporal happening that involves the particles' emission from the source, the selection of measurement settings, concrete measurement processes and detection events, and perhaps more. A satisfactory account of these possible single runs cannot be a mere list of them, and a satisfactory account of an experiment must be more than a mere chronicle of the mentioned characteristics of the actual runs. The challenge is to come up with a law-like account (i.e., one that is in some sense stable under counterfactual variations). Our BST analysis offers a way to represent arguments along these lines in full formal detail.

The main formal tool that we will use is that of a structure extension. Given a BST surface description of a QM correlation experiment—more or less, a compilation of the really possible runs and their spatio-temporal features—we can study ways of extending the surface structure via hidden variables, or instruction sets. The idea is that such an extended structure could show that the modal or probabilistic funny business present in the surface structure is only apparent, and how the surface correlations can come about on the basis of hidden structure without such troublesome correlations.

We begin our work in the setting of modal correlations in Chapter 8.3. We then move to probabilistic correlations in Chapter 8.4. In our exposition, for simplicity's sake we stick to the BST_{92} framework

8.3 Explaining modal correlations via instruction sets

The motivation for wanting to get rid of modal correlations is, bluntly speaking, that we seem unable to understand how they could occur. If two distant events each have a number of different possible outcomes and these events cannot causally influence one another, then how could some joint outcomes be impossible? But that is exactly the situation that we seem to be facing in certain quantum correlation experiments, and which we have called modal funny business (MFB). The task of getting rid of MFB is, therefore, the following: Given a BST structure $\langle W, < \rangle$ that harbors MFB, we ask whether there is a different BST structure $\langle W', <' \rangle$ that is free from MFB while still representing the same facts. This means that we take the initially given BST structure to be a representation of surface facts only, and that we look for an extension of that structure that consists of copies of histories of the surface structure that are differentiated by hidden factors. The idea is that while each history in the surface structure describes the empirically accessible facts about one possible run of the experiment in question, each history in the extended structure describes such a possible run together with hidden factors that are posited to explain the modal correlations present in the surface structure. In agreement with widespread usage, we call these hidden factors instruction sets, or deterministic hidden variables.[10] Technically, the overall idea is to replace one model of a quantum correlation experiment, the surface model, by an extended model that does not contain modal funny business and in which additional instruction sets prescribe the outcomes of measurements (Nature's choices). Such a prescription has to be given for any possible choice of measurement parameters by the experimenter. In general, therefore, an instruction set has to have a counterfactual character; it has to include instructions, not all of which can be realized together.

Indeed, instruction sets without such a counterfactual character appear fishy. As we said earlier and as we will prove below (Theorem 8.2), the following maneuver is always possible. Assume that all the choice points in a BST structure $\langle W, < \rangle$ that models an experiment lie above some event e^*, forming a set E for which $e^* < E$. Take all the histories h that contain e^*—under our assumptions, this is the full set $\mathrm{Hist}(W)$—and create a unique label $\hat{h}$ for each history. Then build a new structure in which each history h, which we can write as $\{a \mid a \in h\}$, is replaced by a labeled copy $h' =_{\mathrm{df}} \{\langle a, \hat{h} \rangle \mid a \in h\}$. Adjust the ordering such that above e^*, the labeled copies are all kept separate, splitting at e^* only. Then the extended structure will have no choice points apart from e^*, so that there can be no modal correlations in the extended structure, but all the histories of the surface structure will still be represented.

Such an extension is formally possible, but it is not considered to be satisfactory as a possible representation of the hidden structure of a quantum experiment. Such extensions are commonly called superdeterministic or conspiratorial, and the reason for this negative verdict is that they

[10] As we said, the BST analysis in this section is special because it proceeds in terms of modal correlations, targeting the underlying structure of possibilities. Usual analyses of hidden variables, on the other hand, take an underlying uncorrelated structure of possibilities for granted and introduce a probability distribution over the hidden factors to explain surface correlations. On that probabilistic approach, a difference is made between *stochastic* hidden variables, which fix probabilities for measurement outcomes, and *deterministic* hidden variables, which *dictate* unique measurement results. For obvious reasons, such deterministic hidden variables are also called instruction sets. See, e.g., Fine (1982, p. 291), who writes of "response functions (giving the λ-determined responses to the measurements)", or Mermin (1981, p. 403), who uses the language of "instruction sets".

eliminate the experimenter's freedom to choose measurement parameters. In a superdeterministic extension, the instruction sets that are selected at e^* not only instruct Nature on the outcomes of the measurements, but they also instruct the experimenter on how to choose the measurement parameters. So the instructions encode a conspiracy between the measurement parameters and the measurement outcomes. This goes against the basic assumption that an experimental outcome is Nature's answer to a question that the experimenter has freely chosen to ask.[11] A satisfactory extension of a surface structure, in contrast, has to retain the experimenter's freedom to choose measurement parameters.

Translated into BST terminology, this means that in the description of the surface structure, we have to distinguish between two kinds of choice points: those that represent Nature's choices of the measurement outcomes (E), and those representing the experimenter's independent choices (C) of experimental parameters. Furthermore, an extension of a surface structure will be satisfactory to the extent to which it eliminates the disturbing MFB among outcomes of members of E while retaining the independence of the choices at members of C.

We will discuss the issue of introducing instruction sets in BST structures in two steps. In Section 8.3.1 we first discuss the general idea of extending a BST surface structure, using the notion of generic instruction sets. Such instruction sets are not fit to be used in real applications, but they allow us to define the general formal procedure of extending a surface structure and to prove a number of important general results about such extensions. For example, we can establish Theorem 8.1, which says that extended structures are again BST_{92} structures. The general discussion paves the way for our discussion of the more specific (and actually useful) notions of non-contextual and contextual instruction sets in Section 8.3.2.

Formally, our point of departure is what we will call a BST_{92} surface structure, in which there is indeterminism induced by experimenters' choices (C) as well as indeterminism produced by Nature (E). The motivation for embarking on the project of a structure extension has to come from some instances of MFB present in the surface structure involving the members of E: some experimental outcomes are modally correlated. (If that is not so, there is no incentive to modify the given surface structure, as it already contains a proper causal account of the possible experimental outcomes.)

[11] See Section 8.1 for some quotes backing this claim.

While removing the MFB present in E, we want to retain the freedom of the experimenters' choices.

In defining a BST_{92} surface structure, the guiding idea is that we are here considering BST as a formal tool for modeling. This means that idealizing assumptions are warranted, as we are only concerned with the experiment in question. We will assume that there are only finitely many choice points in W and that each of them is only finitely splitting. We will also assume that in the model, all choice points are either due to Nature (providing measurement outcomes) or due to the experimenter (providing choices of measurement parameters). Furthermore, the experimenters' choices of measurement parameters must of course be made before the respective measurement with the chosen parameters occurs. Formally, this means that any choice point under the experimenters' control ($c \in C$) must be below some measurement choice point ($e \in E$).

Definition 8.2. A *BST_{92} surface structure* is a quintuple $\langle W, <, e^*, E, C\rangle$, where $\langle W, <\rangle$ is a BST_{92} structure, $e^* \in W$ is a deterministic point in W, and $E, C \subseteq W$ are two finite sets of finitely splitting choice points fulfilling the conditions that $E \neq \emptyset$, $E \cap C = \emptyset$, and that $E \cup C$ is the set of all choice points in W. With respect to the ordering, we demand $e^* < E$ (i.e., for any $e \in E$, we have $e^* < e$), and for any $c \in C$ there must be some $e \in E$ for which $c < e$. For future use we define

$$\tilde{T}_E =_{\mathrm{df}} \{e \rightarrowtail H \mid e \in E, H \in \Pi_e\}; \quad S_E =_{\mathrm{df}} \{T \subseteq \tilde{T}_E \mid T \text{ maximal consistent}\}.$$

Note that $\tilde{T}_E$ and S_E are finite by our assumptions. We will drop the subscript E if it is clear from context.

Here are some simple facts about the structure of the set S of maximal consistent sets of transitions with initials from E.

Fact 8.1. *Let $\langle W, <, e^*, E, C\rangle$ be a BST_{92} surface structure, let $T \in S$, and let $h_1, h_2 \in H(T)$. With respect to the splitting of h_1 and h_2, the following holds: (1) If $h_1 \perp_c h_2$, then $c \in C$. (2) If $h_1 \perp_c h_2$, then there is some $e \in E$ for which $c < e$, but there is no $e \in E$ for which $e \in h_1 \cup h_2$ and for which $c < e$.*

Proof. (1) Let $T_{h_i} =_{\mathrm{df}} \{e \rightarrowtail \Pi_e\langle h_i\rangle \mid e \in E \cap h_i\}$ be the sets of transitions on h_i that have initials in E ($i = 1, 2$). By $h_i \in H(T)$, we have $T \subseteq T_{h_i}$, and as T is a maximal consistent set of transitions with initials in E, in fact $T_{h_1} = T = T_{h_2}$.

So h_1 and h_2 cannot split at a member of E. As $E \cup C$ is the set of choice points in W, it must be that $c \in C$.

(2) By the definition of a BST_{92} surface structure, for any $c \in C$ there must be some $e \in E$ for which $c < e$. By (1), we have $e \in h_1$ iff $e \in h_2$ for members $e \in E$. And if $e \in h_1 \cap h_2$ and $c < e$, then $h_1 \equiv_c h_2$. □

As we said, the idea of adding instruction sets is to explain away surface MFB via a hidden structure that provides instructions at e^* for what should happen at the choice points in E (which represent Nature's choices). Such instructions have to be counterfactual in the sense of allowing for different choices of the experimenters via the choice points in C. In addition, there should be no MFB between the choices of the measurement settings and the measurement outcomes—this would contradict the basic idea that the choice of measurement settings is completely independent of the outcomes that Nature provides. The following definition captures this idea.

Definition 8.3 (C/E independence)**.** We say that a BST_{92} surface structure $\langle W, <, e^*, E, C\rangle$ *violates* C/E *independence* iff there are two consistent, non-empty sets of transitions, T_C, with initials $E_{T_C} \subseteq C$, and T_E, with initials $E_{T_E} \subseteq E$, for which $T_C \cup T_E$ is combinatorially consistent but inconsistent, so that $T_C \cup T_E$ constitutes a case of combinatorial funny business (see Def 5.6). We say that a BST_{92} surface structure (or its set E) *satisfies* C/E *independence* iff it does not violate it.

In the following we develop the idea of instruction sets in a number of ways. In Section 8.3.1 we exhibit the formal procedure of extending a surface structure with respect to an unconstrained notion of instruction sets, which we call *generic* instruction sets. In Section 8.3.2 we then focus on two actually useful types of instruction sets, non-contextual and contextual ones. In Section 8.3.3 we then turn to the analysis, both in terms of non-contextual and contextual instruction sets, of the GHZ experiment, which is a well-known example of modal correlations due to quantum entanglement. We provide a brief summary in Section 8.3.4.

8.3.1 Extensions of a surface structure by generic instruction sets

Generic instruction sets are defined to be, quite simply, subsets of S, that is, sets of maximal consistent sets of transitions with initials from E, the set of Nature's choice points.

Definition 8.4 (Generic instruction set). Let $\langle W, <, e^*, E, C \rangle$ be a BST_{92} surface structure. A *generic instruction set for* $\langle W, <, e^*, E, C \rangle$ is any non-empty subset λ of S, i.e., $\lambda \subseteq S$ and $\lambda \neq \emptyset$. We write $\mathfrak{I}_g$ for the set of all generic instruction sets.

Note that this definition implies, trivially, that any consistent set of transitions can be expanded to a generic instruction set. The notion of a generic instruction set is too wide for applications. The set S generally contains maximal consistent subsets T_1, T_2 that are openly contradictory, prescribing opposite outcomes to each of their initials, and by Def. 8.4, a generic instruction set may contain both T_1 and T_2. In due course, we will exclude such instruction sets from consideration. We work with the notion of a generic instruction set here in order to show that a substantial part of the theory of instruction sets in BST can be developed without any commitment as to the nature of the instruction sets. The first example of such a general definition concerns the notion of matching:

Definition 8.5 (Matching). Let $\lambda \in \mathfrak{I}_g$ be a generic instruction set and h a history in W. We define the *matching set for h* to be $T_h =_{\mathrm{df}} \{e \rightarrowtail \Pi_e\langle h \rangle \mid e \in E \cap h\}$. We say that *h matches* λ iff λ contains the matching set for h, i.e., iff $T_h \subseteq \bigcup \lambda$. As a stylistic variant, we also say that λ *matches h*.

By this definition, if a history h contains no elements of E, then $T_h = \emptyset$, and *any* instruction set $\lambda \in \mathfrak{I}_g$ matches h. For the general case, we can also show that there are always matching pairs of histories and instruction sets.

Fact 8.2. *(1) For any history h in W there is some generic instruction set $\lambda \in \mathfrak{I}_g$ that matches h. (2) Let $\lambda \in \mathfrak{I}_g$ and $T \in \lambda$, and let h be a history. If $h \in H(T)$, then h matches λ. (3) For any generic instruction set $\lambda \in \mathfrak{I}_g$, there is some matching history h.*

Proof. (1) Let $h \in \mathrm{Hist}(W)$ be given. The matching set for h, T_h of Def. 8.5, is consistent, so it can be extended to some maximal consistent set $T \in S$. The set $\lambda =_{\mathrm{df}} \{T\}$ is already a generic instruction set according to Def. 8.4. The instruction set λ (and any of its extensions) matches h because $T_h \subseteq \bigcup \lambda$.

(2) Let $\lambda \in \mathfrak{I}_g$ and $T \in \lambda$ be given. The set T is consistent. Assume that $h \in H(T)$, and consider the matching set for h, T_h of Def. 8.5. As h lies in all outcomes of transitions of T, which all have initials in E, it must be that $T \subseteq T_h$. Now T is a maximal consistent subset of $\tilde{T}$, so T_h, which is also consistent, cannot be a proper superset of T, so that $T = T_h$. From $T \in \lambda$ we have $T \subseteq \bigcup \lambda$, so that indeed $T_h \subseteq \bigcup \lambda$, i.e., h matches λ.

(3) As $\lambda \neq \emptyset$, there is some $T \in \lambda$, which is consistent. Take $h \in H(T)$; the claim follows directly from (2). □

If a history h in a BST_{92} surface structure matches an instruction set λ, this means that λ provides a direction (a possible outcome) at each member $e \in E$ that occurs on h, and that direction is "to stay on h" ($\Pi_e\langle h\rangle$). The instruction set λ is, as it were, h-friendly. Given the generality of the definition of a generic instruction set, however, the set of transitions $\bigcup\lambda$ may in such a case also contain a different, incompatible transition with initial e, $e \rightarrowtail H$ with $H \neq \Pi_e\langle h\rangle$. So, on the generic approach, an instruction set can match a history even though the instructions do not *dictate* that the history occur—they only have to *allow* for the history to occur. This feature of the possibility of different instructions for the same initial is what distinguishes non-contextual from contextual instruction sets: for the former, an instruction at one of Nature's choice points $e \in E$ has to be unique, while for the latter, it may depend on the context and thus fail to be unique.

Given the notion of matching, we can work toward our definition of an extended structure (Def. 8.7). In an extended structure, we replace elements $a \in W$ of the surface structure with labeled elements $\langle a, L\rangle$. The labels represent instruction sets, which will be implemented as new elementary outcomes of e^*. So for $a > e^*$, the label L has to be some $\lambda \in \mathfrak{I}_g$. For $a \not> e^*$, on the other hand, an outcome at e^* can have no causal influence on a; in this case we use the label $L = \emptyset$, just in order to preserve a uniform format for the members of the extended structure. Formally, we build the extended structure from lifted histories, which are defined as follows:

Definition 8.6 (Lifted history)**.** Let h be a history in W matching a generic instruction set λ. Then we define the lifted history $\varphi_\lambda(h)$ to be

$$\varphi_\lambda(h) =_{\text{df}} \{\langle a, \emptyset\rangle \mid a \in h, a \not> e^*\} \cup \{\langle a, \lambda\rangle \mid a \in h, a > e^*\}.$$

Using these lifted histories, we define the *extended BST structure* based on a BST_{92} surface structure as follows:

Definition 8.7 (Extended structure)**.** Let $\mathscr{W}_S = \langle W, <, e^*, E, C\rangle$ be a BST_{92} surface structure, and let $\mathfrak{I}_g$ be its set of generic instruction sets. We define the extended structure $\mathscr{W}_E = \langle W', <'\rangle$ corresponding to $\mathscr{W}_S$ to be the union of all lifted histories together with an ordering relation that respects that events can be multiply copied with different instruction sets.

$$W' =_{\mathrm{df}} \bigcup_{h \in \mathrm{Hist}(W),\ \lambda \in \mathfrak{I}_g \text{ matching } h} \varphi_\lambda(h);$$

$$\langle e_1, L_1 \rangle <' \langle e_2, L_2 \rangle \text{ iff } e_1 < e_2 \text{ and } (L_1 = \emptyset \text{ or } L_1 = L_2);$$

$$\mathscr{W}_E =_{\mathrm{df}} \langle W', <' \rangle.$$

The different copies of elements above e^* are not order related. The following fact about elements of an extended structure shows how elements of W' come from lifted histories:

Fact 8.3. *Let $\mathscr{W}_E = \langle W', <' \rangle$ be the extended structure based on a BST$_{92}$ surface structure $\mathscr{W}_S = \langle W, <, e^*, E, C \rangle$. Then $\langle a, L \rangle \in W'$ iff either ($a \not> e^*$ and $L = \emptyset$) or ($a > e^*$ and there is some history h matching λ for which $a \in h$).*

Proof. For $a \not> e^*$, the "⇒" direction follows by Def. 8.7. For the "⇐" direction, note that a belongs to some history $h \in \mathrm{Hist}(W)$, and by Fact 8.2(1), there is some $\lambda \in \mathfrak{I}_g$ matching h.

For $a > e^*$, the claim follows immediately from Def. 8.7. □

We are working toward our first main result about extended structures, Theorem 8.1, which says that these are also BST$_{92}$ structures. As a simple first step, we can show that some basic properties of the ordering $<$ carry over to $<'$ immediately:

Fact 8.4. *Let $\mathscr{W}_E = \langle W', <' \rangle$ be the extended structure based on a BST$_{92}$ surface structure $\mathscr{W}_S = \langle W, <, e^*, E, C \rangle$. Then $\langle W', <' \rangle$ is a non-empty, dense, strict partial ordering.*

Proof. Left as Exercise 8.1 □

We now have to characterize the histories in the extended structure, so that we can prove further BST$_{92}$-relevant properties. We split the crucial history lemma (Lemma 8.1) into two facts. The one direction is the following.

Fact 8.5. *Let $\langle W, <, e^*, E, C \rangle$ be a BST$_{92}$ surface structure, let $h \in \mathrm{Hist}(W)$, and let $\lambda \in \mathfrak{I}_g$ be a generic instruction set matching h. Then the set $\varphi_\lambda(h)$ is maximal directed, i.e., it is a history in the extended structure $\langle W', <' \rangle$.*

Proof. Let $A' =_{\mathrm{df}} \varphi_\lambda(h)$ be the lifted history. We first show that A' is directed: Let $\langle e_1, L_1 \rangle, \langle e_2, L_2 \rangle \in A'$ (where $L_i = \lambda$ or $\emptyset$ depending on whether or not $e_i > e^*$, $i = 1, 2$). These elements of A' were lifted from elements $e_1, e_2 \in h$, and as h is directed, there is some $e_3 \in h$ for which $e_1 \leqslant e_3$

and $e_2 \leqslant e_3$. The corresponding element $\langle e_3, L_3 \rangle$ of A' is $\leqslant'$-above the two mentioned elements of A', as one can easily verify by the definition of the ordering.

Assume now for reductio that A' is not maximal directed, i.e., that there is a proper superset $A'' \supsetneq A'$, $A'' \subseteq W'$, that is also directed. As a subset of W', A'' has elements $\langle a, L \rangle$, where $L = \emptyset$ iff $a \not> e^*$. We define the set $A =_{\mathrm{df}} \{a \in W \mid \text{there is some } \langle a, L \rangle \in A''\}$. We consider two cases.

Case 1: If A' contains some element $\langle a_0, \lambda \rangle$, then any $\langle a, L \rangle \in A''$ for which $a > e^*$ must satisfy $L = \lambda$; this follows by directedness of A'' and by the definition of the ordering. So if $\langle a, L \rangle, \langle a, L' \rangle \in A''$, then $L = L'$; the labels of elements of A'' are unique. We can show that the set $A \subseteq W$ defined above is directed: Pick $e_1, e_2 \in A$, so that the respective $\langle e_1, L_1 \rangle, \langle e_2, L_2 \rangle$ are members of A''. As A'' is directed, these two members have a common upper bound $\langle e_3, L_3 \rangle \in A''$, so that there is $e_3 \in A$ that is a common upper bound of e_1 and e_2 in W. Now there is some $\langle e_0, L_0 \rangle \in A'' \setminus A'$ by our reductio assumption, so that A contains some member $e_0 \notin h$. So A is a proper superset of h that is also directed, which contradicts the definition of histories as maximal directed sets.

Case 2: If A' contains only elements that have the label $\emptyset$, the directed proper superset A'' (whose existence constitutes our reductio assumption) might also contain only elements that have the label $\emptyset$. So the labels of elements of A are unique in A'', and we can reason as in case 1. It might be, however, that A'' contains an element $\langle a_0, \lambda' \rangle$ with $\lambda' \neq \emptyset$. In this case we cannot guarantee that $\lambda' = \lambda$. But it must still be that any $\langle a, L \rangle \in A''$ for which $a > e^*$ must satisfy $L = \lambda'$, again by directedness and by the definition of the ordering. So in this case too, the labels of elements of A are unique, and we can again reason exactly as in case 1. □

For the second direction of the history lemma, we have to use the assumption that the indeterminism in W is finite.[12]

Fact 8.6. *Let $\langle W, <, e^*, E, C \rangle$ be a BST_{92} surface structure. Let $A' \subseteq W'$ be maximal directed. Then there is some $h \in \mathrm{Hist}(W)$ and some λ matching h such that $A' = \varphi_\lambda(h)$.*

[12] It is instructive to see how infinite structures can cause trouble here. If, for example, there is an infinite chain of choice points from C that has no maximum in the set called A in the proof below, it may be impossible to find a history containing all of A that matches a given instruction set λ.

Proof. Let $A =_{\text{df}} \{a \in W \mid \langle a, L\rangle \in A'\}$. As in the proof of Fact 8.5, using the definition of the ordering $<'$ we can establish that A is directed.

We show that (*) there is some $a^* \in A$ for which any $h^* \in H_{a^*}$ contains all of A. Consider the set of past cause-like loci for members of A,

$$B =_{\text{df}} \bigcup_{a \in A} \{c \in W \mid c < a \text{ and there is some } h \text{ for which } h \perp_c H_a\}.$$

Note that B contains only indeterministic events, so by the finiteness of indeterminism in W, B is finite of size N. For each of the finitely many $c_i \in B$, we can pick some $a_i \in A$ for which $c_i < a_i$ ($i = 1, \ldots, N$). At this point we also tend to a subtlety concerning the labels of elements of A', as our sought-after element $a^* \in A$ must admit a label that is appropriate for all of A. So if A' contains some element $\langle a, L\rangle$ with $L \neq \emptyset$, then this $L \in \mathfrak{I}_g$ must be unique in the directed set A' by the definition of the ordering, and so we let $a_0 =_{\text{df}} a$. Otherwise, all elements of A' have the label $\emptyset$, and a_0 is not needed; for a uniform construction, we simply set $a_0 =_{\text{df}} a_1$. The element a_0 thus keeps track of which labels occur in A'. As A is directed, there is some $a^* \in A$ for which all the finitely many $a_i < a^*$ ($i = 0, \ldots, N$), and thereby $c_i < a^*$ ($i = 1, \ldots, N$). Now pick some $h_A \in H_{[A]}$; such a history exists as A is directed. Consider an arbitrary $h^* \in H_{a^*}$. Note that by the choice of a^*, for any $c_i \in B$ we have $h^* \equiv_{c_i} h_A$ because $a_i \in h^* \cap h_A$ and $c_i < a_i$. We claim that $A \subseteq h^*$. Assume not, then there must be some $a \in A \setminus h^*$, whence $a \in h_A \setminus h^*$. By PCP, there must then be some $c < a$ for which $h^* \perp_c h_A$ and in fact $h^* \perp_c H_a$, so by the definition of B, $c \in B$. This implies that $c = c_i$ for some $i \in \{1, \ldots, N\}$, but we have established $h^* \equiv_{c_i} h_A$, contradicting $h^* \perp_c h_A$. So indeed, any history h^* that contains a^* contains all of A. Now as $a^* \in A$, we have $\langle a^*, L\rangle \in A'$ with $L = \emptyset$ or $L = \lambda$ for some $\lambda \in \mathfrak{I}_g$. Given the way the element a_0 was picked, in the former case all elements of A must have the label $\emptyset$ in A'; pick some $h^* \in H_{a^*}$. By Fact 8.2(1), we can pick some λ that matches h^*. In the latter case, as $\langle a^*, \lambda\rangle \in W'$, by Fact 8.3 there is some $h^* \in H_{a^*}$ that matches λ. In both cases, we have a history $h^* \in H_{a^*}$ that matches λ, and by (*), that h^* contains all of A.

So by Fact 8.5, the set $A^* =_{\text{df}} \varphi_\lambda(h^*)$ is a maximal directed subset of W', and as $A \subseteq h^*$, we have that A^* is a superset of A'. Now as A' is maximal directed, it has to be that $A^* = A'$, i.e., $A' = \varphi_\lambda(h^*)$. □

Given the two above Facts, we have established our history lemma:

Lemma 8.1. *Let $\langle W, <, e^*, E, C\rangle$ be a BST_{92} surface structure. The set $A' \subseteq W'$ is a history in W' (maximal directed) if and only if there is some $h \in \mathrm{Hist}(W)$ and some generic instruction set λ matching h such that $A' = \varphi_\lambda(h)$.*

Proof. The two directions have been shown as Facts 8.5 and Facts 8.6. □

Now we can go on to show that extended structures are in fact BST_{92} structures:

Theorem 8.1. *Let $\langle W, <, e^*, E, C\rangle$ be a BST_{92} surface structure, and let $\langle W', <'\rangle$ be the corresponding extended structure. Then $\langle W', <'\rangle$ is a BST_{92} structure.*

Proof. By Fact 8.4, $\langle W', <'\rangle$ is a dense, strict partial ordering. It remains to prove that there are infima for lower bounded chains and history-relative suprema for upper bounded chains, that Weiner's postulate holds, and that the prior choice postulate, PCP_{92}, is satisfied. Given Lemma 8.1, all these properties of $\langle W', <'\rangle$ can be lifted from the respective properties of $\langle W, <\rangle$. We ask the reader to supply details via Exercise 8.2. □

Note that in the whole series of proofs leading up to Theorem 8.1, apart from being a subset of S, the inner structure of the instruction sets has played no role. This means that we are free to work with more restrictive notions of instruction sets in real applications without having to revisit the whole construction.

8.3.1.1 The possibility of superdeterministic extensions

While motivating the distinction between Nature's choices at $e \in E$ and experimenters' choices at $c \in C$, we pointed to the possibility of superdeterministic extensions. We are now in a position to define them in more formal detail.

The basic idea of a superdeterministic extension of a surface structure is that *all* of the choice points in the structure are taken care of via the instruction sets at e^*. A single choice at e^* determines all the outcomes of all the choice points in its future. Formally speaking, we have such a situation in case $C = \emptyset$ and the event e^* is in the common past of all the choice points ($e^* < E$). One might say that this amounts to taking the experimenters to be part of Nature, so that the instruction sets pertain to their actions as well as to the experimental outcomes. As we remarked, this move is generally not considered to be satisfactory.[13] Here we are only concerned with spelling out what the move amounts to in the BST framework.[14]

In case $C = \emptyset$, the set $\tilde{T}$ of transitions from members of E is the total set of indeterministic transitions in W, $\mathrm{TR}(W) = \tilde{T}$. Accordingly, any member $T \in S$ singles out exactly one history from $\mathrm{Hist}(W)$, which we denote h_T; that history's matching set is again T itself:

$$\text{for } T \in S, \quad H(T) = \{h_T\} \quad \text{and} \quad T_{h_T} = T.$$

Accordingly, any $T \in S$ matches exactly one history, viz., h_T. In this case, each appropriate instruction set should single out exactly one history. This amounts to taking superdeterministic instruction sets to be singletons of members of S:

Definition 8.8 (Superdeterministic instruction set)**.** Let $\langle W, <, e^*, E, C\rangle$ be a BST$_{92}$ surface structure in which $C = \emptyset$. A *superdeterministic instruction set for* $\langle W, <, e^*, E, C\rangle$ is a singleton $\lambda = \{T\}$ with $T \in S$. We write $\mathfrak{I}_s$ for the set of all superdeterministic instruction sets.

As we have already remarked, all of our above results about structure extensions stay in place for restrictions of generic instruction sets, including superdeterministic instruction sets. So, for a BST$_{92}$ surface structure $\langle W, <, e^*, E, C\rangle$ with $C = \emptyset$, the *superdeterministic extension* is well-defined:

Definition 8.9. Let $\langle W, <, e^*, E, C\rangle$ be a BST$_{92}$ surface structure in which $C = \emptyset$, and let $\mathfrak{I}_s$ be its set of superdeterministic instruction sets. The corresponding *superdeterministic extension* is the extended structure $\langle W', <'\rangle$ of Def. 8.7, replacing $\mathfrak{I}_g$ by $\mathfrak{I}_s$.

Given these definitions, we can establish our main theorem about superdeterministic extensions:

Theorem 8.2. *Let $\langle W, <, e^*, E, C\rangle$ be a BST$_{92}$ surface structure in which $C = \emptyset$, and let $\langle W', <'\rangle$ be its corresponding superdeterministic extension. Then $\langle W', <'\rangle$ is a BST$_{92}$ structure without MFB in which there is exactly one choice point, $\langle e^*, \emptyset\rangle$. Furthermore, there is a bijection between the sets of histories* $\mathrm{Hist}(W)$ *and* $\mathrm{Hist}(W')$.

[13] For an interesting dissenting voice, see Adlam (2018), who combines a notion of global determinism with a reassessment of quantum correlations from a perspectivalist point of view.

[14] We are convinced that our formal analysis provides useful material for a discussion of the status of experiments in the free will debate, but we leave this matter to one side here.

Proof. As we noted, $\langle W', <' \rangle$ is a BST_{92} structure by Theorem 8.1. Any history $h \in \text{Hist}(W)$ matches exactly one instruction set $\lambda_h =_{\text{df}} \{T_h\}$, so by Lemma 8.1, the mapping

$$h \mapsto \varphi_{\lambda_h}(h)$$

is a bijection between $\text{Hist}(W)$ and $\text{Hist}(W')$. Note that $E \neq \emptyset$, so the set S has at least two members. Accordingly, W and W' have at least two histories. Let $h_1', h_2' \in \text{Hist}(W')$ with $h_1' \neq h_2'$; by the previous observation, we must have $h_i' = \varphi_{\lambda_i}(h_i)$ for $h_i \in \text{Hist}(W)$ and $\lambda_i = \{T_{h_i}\}$ $(i = 1, 2)$, $\lambda_1 \neq \lambda_2$. By the definition of the ordering, the point $\langle e^*, \emptyset \rangle$ is maximal in $h_1' \cap h_2'$; that is, a choice point for the arbitrary histories $h_1', h_2' \in \text{Hist}(W')$. As W' contains only one choice point, there can be no MFB in W'. □

The upshot of this result is simple: if a quantum correlation experiment is modeled in such a way that instruction sets (deterministic hidden variables) are allowed to determine not just measurement outcomes but also all of the experimenters' choices ($C = \emptyset$), then one can replace the surface structure, no matter how many cases of modal funny business it contains, with a very simple structure that represents the same surface facts (has exactly isomorphic histories) while pulling together all indeterminism to a single point in the past of E. This result is not surprising, but it is good as a reality check for our formal approach before we embark on a more general discussion of splittings induced by instruction sets.

8.3.1.2 Splitting in extended structures: The general case

In the general case, $C \neq \emptyset$, and we need to characterize the splitting of histories in an extended structure in relation to the splitting of histories in the initially given surface structure. Our aim was that an extended structure should be instructed to behave in some specific way toward the members of E (Nature's choices should be guided by the instruction set: ideally, Nature will have no choice at all apart from the splitting at e^*), whereas the indeterminsm outside of E should be retained (experimenters should still be free to choose measurement settings at choice points $c \in C$, no matter which instruction set for Nature has been given at e^*). The lemma below shows to what extent these aims can be accomplished for *generic* instruction sets. This will also provide helpful guidance toward the construction of practically useful notions of instruction sets later on.

Lemma 8.2. *We consider a pair of histories in W and a corresponding pair of histories in W', as provided by Lemma 8.1: Let $h_1, h_2 \in \mathrm{Hist}(W)$, let $\lambda_1, \lambda_2 \in \mathfrak{I}_g$ be instruction sets, let λ_i match h_i, and set $h'_i =_{\mathrm{df}} \varphi_{\lambda_i}(h_i)$ ($i = 1,2$). The splitting of histories in W and in W' can be characterized as follows:*

(1) As to splitting at e^*, we have $h_1 \equiv_{e^*} h_2$, and we have $h_1 \equiv_{\langle e^*,\emptyset\rangle} h_2$ iff $\lambda_1 = \lambda_2$.
(2) Let $c \not> e^*$ and $c \neq e^*$ (which implies $c \notin E$). We have $h_1 \equiv_c h_2$ iff $h'_1 \equiv_{\langle c,\emptyset\rangle} h'_2$.
(3) Let $c \notin E$ and $c > e^*$ and assume $\lambda_1 = \lambda_2 = \lambda$. Then $h_1 \equiv_c h_2$ iff $h'_1 \equiv_{\langle c,\lambda\rangle} h'_2$.
(4) Let $c \in E$ (hence $c > e^*$) and $\lambda_1 \neq \lambda_2$. If $h_1 \perp_c h_2$, then neither $h'_1 \perp_{\langle c,\lambda_1\rangle} h'_2$ nor $h'_1 \perp_{\langle c,\lambda_2\rangle} h'_2$, but $h'_1 \perp_{\langle e^*,\emptyset\rangle} h'_2$.
(5) Let $c \in E$, $\lambda_1 = \lambda_2 = \lambda$, and assume that $\bigcup \lambda$ contains no blatantly inconsistent transitions with initial c. Then it cannot be that $h_1 \perp_c h_2$, nor that $h'_1 \perp_{\langle c,\lambda\rangle} h'_2$.
(6) Let $c \in E$, $\lambda_1 = \lambda_2 = \lambda$, and assume that $\bigcup \lambda$ contains two blatantly inconsistent transitions $c \rightarrowtail \Pi_c\langle h_1\rangle \in T_1$, $c \rightarrowtail \Pi_c\langle h_2\rangle \in T_2$, with $T_1, T_2 \in \lambda$. Then $h_1 \perp_c h_2$, and $h'_1 \perp_{\langle c,\lambda\rangle} h'_2$.
(7) Let $c \in E$, $\lambda_1 = \lambda_2 = \lambda$, and $h_1 \equiv_c h_2$. Then $h'_1 \equiv_{\langle c,\lambda\rangle} h'_2$.

Proof. (1) By definition e^* is deterministic, so $h_1 \equiv_{e^*} h_2$. And by the definition of the ordering, if $\lambda_1 \neq \lambda_2$, there is no $e > e^*$ and no label L for which $\langle e, L\rangle \in h'_1 \cap h'_2$. On the other hand, $h_1 \equiv_{e^*} h_2$ implies that there is $e \in h_1 \cap h_2$ with $e^* < e$. So if $\lambda_1 = \lambda_2 = \lambda$, we have $\langle e, \lambda\rangle \in h'_1 \cap h'_2$, and hence $h'_1 \equiv_{\langle e^*,\emptyset\rangle} h'_2$.

(2) Let $c \not> e^*$, so $c \in h_i$ iff $\langle c, \emptyset\rangle \in h'_i$ ($i = 1,2$). If $c \notin h_1 \cap h_2$, both undividedness claims are trivially false, so assume $c \in h_1 \cap h_2$, which implies that we have $\langle c, \emptyset\rangle \in h'_1 \cap h'_2$ as well. We have to show that $h_1 \equiv_c h_2$ iff $h'_1 \equiv_{\langle c,\emptyset\rangle} h'_2$.
"$\Rightarrow$": Assume $h_1 \equiv_c h_2$ but $h'_1 \not\equiv_{\langle c,\emptyset\rangle} h'_2$, which by the above implies $h'_1 \perp_{\langle c,\emptyset\rangle} h'_2$. By $h_1 \equiv_c h_2$, there is $a \in h_1 \cap h_2$ above c. If $a \not> e^*$, then $\langle a, \emptyset\rangle \in h'_1 \cap h'_2$ and $\langle c, \emptyset\rangle <' \langle a, \emptyset\rangle$, contradicting $h'_1 \perp_{\langle c,\emptyset\rangle} h'_2$. If $a > e^*$, we will find some $a' > c$ for which $a' \not> e^*$. To this end consider a maximal chain l containing c and a. Then $l' =_{\mathrm{df}} \{x \in l \mid x > e^*\}$ is a non-empty chain (it contains a) that is lower bounded by c, so it has an infimum i. By density, there is some a' for which $c < a' < i$, and $a' \not> e^*$. Thus, $\langle a', \emptyset\rangle \in h'_1 \cap h'_2$ and $\langle c, \emptyset\rangle <' \langle a', \emptyset\rangle$, contradicting $h'_1 \perp_{\langle c,\emptyset\rangle} h'_2$.

"$\Leftarrow$": Assume $h'_1 \equiv_{\langle c,\emptyset\rangle} h'_2$ but $h_1 \not\equiv_c h_2$. The latter implies $h_1 \perp_c h_2$. The former must be witnessed by some $\langle a,L\rangle \in h'_1 \cap h'_2$, where $\langle c,\emptyset\rangle <' \langle a,L\rangle$. Hence, by the definition of the ordering, $a \in h_1 \cap h_2$ and $c < a$, which contradicts $h_1 \perp_c h_2$.

(3) This is shown exactly as (2). We show that if $c \notin h_1 \cap h_2$, then the equivalence holds trivially (each side is false). Next we show that $c \in h_1 \cap h_2$ iff $\langle c,\lambda\rangle \in h'_1 \cap h'_2$. Finally, we run two reductio arguments: (1) assume $h_1 \equiv_c h_2$ but $h'_1 \not\equiv_{\langle c,\emptyset\rangle} h'_2$ and (2) assume $h'_1 \equiv_{\langle c,\emptyset\rangle} h'_2$ but $h_1 \not\equiv_c h_2$. Both assumptions lead to a contradiction, as in (2).

(4) Let $c > e^*$ and $\lambda_1 \neq \lambda_2$. Assume $h_1 \perp_c h_2$. Note that h'_1 and h'_2 contain different copies of c, $\langle c,\lambda_1\rangle \neq \langle c,\lambda_2\rangle$, so the presupposition for being undivided ($h'_1 \equiv_{\langle c,\lambda_i\rangle} h'_2$) and for splitting ($h'_1 \perp_{\langle c,\lambda_i\rangle} h'_2$) is violated, so none of these can hold ($i = 1,2$). In this case, by the definition of the ordering, $\langle e^*,\emptyset\rangle$ is maximal in the intersection of h'_1 and h'_2, i.e., $h'_1 \perp_{\langle e^*,\emptyset\rangle} h'_2$.

(5) Assume for reductio that $h_1 \perp_c h_2$. This implies that $\tau_1 =_{\mathrm{df}} (c \rightarrowtail \Pi_c\langle h_1\rangle) \neq \tau_2 =_{\mathrm{df}} (c \rightarrowtail \Pi_c\langle h_2\rangle)$. But as h_1 and h_2 both match λ by assumption and $c \in h_1 \cap h_2$, it has to be that $\tau_1 \in \bigcup\lambda$ (by matching h_1) and $\tau_2 \in \bigcup\lambda$ (by matching h_2), which contradicts the assumption that $\bigcup\lambda$ contains no blatantly inconsistent transitions with initial c. The assumption that $h'_1 \perp_{\langle c,\lambda\rangle} h'_2$ can be dealt with via that former case: It cannot be that $h_1 \cap h_2$ contains some $a > c$, for then $\langle a,\lambda\rangle > \langle c,\lambda\rangle$, showing $h'_1 \equiv_{\langle c,\lambda\rangle} h'_2$. So $h_1 \perp_c h_2$, and we continue as above.

(6) As λ matches both h_1 and h_2, the existence of such blatantly inconsistent transitions implies $h_1 \perp_c h_2$. Thereby we have $\langle c,\lambda\rangle \in h'_1 \cap h'_2$. Now it cannot be that $h'_1 \equiv_{\langle c,\lambda\rangle} h'_2$: this would imply that there is $\langle a,\lambda\rangle \in h'_1 \cap h'_2$ for which $\langle a,\lambda\rangle > \langle c,\lambda\rangle$, so that there is $a \in h_1 \cap h_2$ with $a > c$, contradicting $h_1 \perp_c h_2$. Thus, $h'_1 \perp_{\langle c,\lambda\rangle} h'_2$.

(7) As $h_1 \equiv_c h_2$, there is $c_1 > c$ such that $c_1 \in h_1 \cap h_2$. Hence $\langle c,\lambda\rangle < \langle c_1,\lambda\rangle$ and $\langle c_1,\lambda\rangle \in h'_1 \cap h'_2$. Thus, $h'_1 \equiv_{\langle c,\lambda\rangle} h'_2$. □

Let us summarize these results in plain English. The biggest change that a structure extension brings concerns e^*. The event e^* is deterministic in the original structure by assumption, but its counterpart $\langle e^*,\emptyset\rangle$ is a new seed of Nature's indeterminism in the extended structure, with its elementary outcomes playing the role of instruction sets (clause 1). Next, there are no changes with respect to choices in the region not above e^* (excluding e^* itself), by clause (2). As noted, such choices cannot involve members of E, as $e^* < E$. Together with clause (3), clause (2) implies that choices outside of E are preserved with respect to all appropriate instruction sets λ. Note that counterparts of clauses (2) and (3) also hold for the splitting relations $\perp_c$ and $\perp_{\langle c,L\rangle}$. Clause (4) shows that the structure extension leads to different copies of members of E for different instruction sets, placing the respective splitting at $\langle e^*, \emptyset\rangle$. Clauses (5) and (6) point to an important distinction that we will spell out below, when we specify two useful notions of instruction sets, viz., contextual vs. non-contextual ones. The difference between them lies exactly in whether $\bigcup\lambda$ is allowed to contain blatantly inconsistent pairs of transitions (contextuality) or not (non-contextuality). Clauses (5) and (6) inform about the consequences. They say that in the absence of contextuality (in the absence of blatantly inconsistent pairs of transitions with a given initial c in $\bigcup\lambda$), Nature's choice at $c \in E$ is completely removed and replaced by the new splitting at $\langle e^*, \emptyset\rangle$ (clause (5)), while splittings at a member $c \in E$ are retained in case an instruction set λ does not give a unique verdict for what has to happen at c.

Earlier, we defined the notion of C/E-independence for BST$_{92}$ surface structures (Def. 8.3). The guiding idea was that experimenters' choices (outcomes of choice points $c \in C$) should be independent of the outcomes of Nature's choices at choice points $e \in E$. A BST$_{92}$ surface structure is C/E-independent iff there is no modal funny business involving outcomes of both members of C and members of E. With respect to the splitting in the extended structure, a similar question can be asked: are there modal correlations involving Nature's choice of an instruction set at the new splitting point $\langle e^*, \emptyset\rangle$, or Nature's remaining choices above e^*, and sets of experimenters' choices? The following definition provides the relevant notion of C/Ext-independence.[15]

Definition 8.10 (C/Ext independence)**.** We say that an extended structure $\langle W', <'\rangle$ derived from a BST$_{92}$ surface structure $\langle W, <, e^*, E, C\rangle$ *violates C/Ext independence* iff there is a case of MFB that involves some C-based and some E-based transitions in one of the following two ways:

(1) There is some $\lambda \in \mathfrak{I}_g$ and a transition

$$\tau^{\lambda}_{e^*} = \langle e^*, \emptyset\rangle \rightarrowtail \Pi_{\langle e^*,\emptyset\rangle}\langle \varphi_\lambda(h)\rangle, \quad \text{with } h \in H_{e^*},$$

and a consistent, non-empty set of transitions T'_C with initials $\langle c, L\rangle$, $c \in C_0 \subseteq C$ and $L \in \{\emptyset, \lambda\}$, for which $T'_C \cup \{\tau^{\lambda}_{e^*}\}$ is combinatorially consistent but

[15] As with other definitions and results in this general part, we write it out using generic instructions sets for concreteness, but the definition is exactly the same for other types of instruction sets.

inconsistent, thus constituting a case of combinatorial funny business (see Def 5.6).

(2) There is some $\lambda \in \mathfrak{I}_g$, a consistent, non-empty set of transitions T'_E with initials $\langle e, L\rangle$, $e \in E_0 \subseteq E$ and $L \in \{\emptyset, \lambda\}$, and a consistent, non-empty set of transitions T'_C with initials $\langle c, L\rangle$, $c \in C_0 \subseteq C$ and $L \in \{\emptyset, \lambda\}$, for which $T'_C \cup T'_E$ constitutes a case of combinatorial funny business.

We say that the extension *satisfies C/Ext independence* iff it does not violate it.

Here we cannot yet prove any general results about C/Ext-independence, apart from the following triviality:

Fact 8.7. *A superdeterministic extension is C/Ext independent.*

Proof. The only choice point in such an extension is $\langle e^*, \emptyset\rangle$, so there can be no MFB in a superdeterministic extension. □

8.3.2 Non-contextual and contextual instruction sets

So far, we have discussed two extreme cases of structure extensions: generic ones ($\mathfrak{I}_g$), in which there are no constraints on instruction sets at all, and superdeterministic ones ($\mathfrak{I}_s$), in which instructions pertain to *all* choice points in the surface structure, so that an instruction set given out at the new splitting point $\langle e^*, \emptyset\rangle$ amounts to the selection of a single history. We already pointed out that these two extreme cases are not satisfactory from a philosophical point of view. The challenge for a useful notion of instruction sets is to steer between these two extremes and to offer instruction sets that are comprehensive enough to allow for the free choice of experimental parameters while still providing useful guidance for the outcomes of Nature's choice points. In terms of structure extensions, this translates into two demands: (1) to allow for surface structures in which $C \neq \emptyset$ (unlike in superdeterminism) and (2) to provide a determinate outcome at the members of E that occur, given the instruction set λ.

Demand (2) is still a bit vague, and there is a good reason for that, as there appear to be two ways to fulfill it: the already mentioned non-contextual and contextual approaches to instruction sets. In line with the previous notation, we denote the sets of these instruction sets, which will again be sets of subsets of S, as $\mathfrak{I}_n$ and $\mathfrak{I}_c$, respectively. It is easiest to explain non-contextual instruction sets first, as these are more tightly constrained than the properly contextual ones. A non-contextual instruction set λ provides instructions for all or for a part of Nature's choice points $e \in E$, specifying exactly one outcome of e for a transition $\lambda(e) = e \rightarrowtail H$, $H \in \Pi_e$. A contextual instruction set, on the other hand, may specify different outcomes for one and the same choice point $e \in E$, but not arbitrarily: if there is a difference in outcome, there must also be a difference in the measurement context (hence the name).

8.3.2.1 Non-contextual instruction sets

We first define non-contextual instruction sets.

Definition 8.11 (Non-contextual instruction sets). Given a BST_{92} surface structure $\langle W, <, e^*, E, C\rangle$, the set of sets of transitions $\lambda \subseteq S$ is a *non-contextual instruction set for* $\langle W, <, e^*, E, C\rangle$ iff λ is maximal with respect to the conditions that (1) $\bigcup\lambda$ is not blatantly inconsistent and (2) for every consistent set of initials of transitions in $\bigcup\lambda$, the respective set of transitions (which is uniquely determined by (1)) is also consistent. We write $\mathfrak{I}_n$ for the set of all non-contextual instruction sets.

To unpack this definition, by condition (1), any λ can be viewed as a partial function from E into $\tilde{T}$, providing exactly one transition with initial e for every $e \in E_\lambda$, where $E_\lambda \subseteq E$ is the set of initials of transitions in $\bigcup\lambda$. We will write $\lambda(e) = \tau$ to indicate that τ is the unique transition with the initial e that occurs in $\bigcup\lambda$, and we extend this notation to sets of initials, so that for $E_0 \subseteq E_\lambda$, $\lambda(E_0) = T_0$ means that $T_0 \subseteq \bigcup\lambda$ is the unique set of transitions with initials in E_0 that occur in λ. So condition (2) can be written as follows: if $E_0 \subseteq E_\lambda$ is consistent, then $\lambda(E_0) \subseteq \tilde{T}$ is also consistent.

We can characterize condition (1) of Def. 8.11 also in a different way, which paves the way for a generalization to contextual instruction sets. The relevant fact is this:

Fact 8.8. *For $\lambda \subseteq S$, define $H(\lambda) = \bigcup_{T\in\lambda} H(T)$. For a set $\lambda \subseteq S$, $\bigcup\lambda$ is blatantly inconsistent iff there are $h_1, h_2 \in H(\lambda)$ for which $h_1 \perp_e h_2$ for some $e \in E$. Accordingly, $\bigcup\lambda$ is not blatantly inconsistent iff for all $h_1, h_2 \in H(\lambda)$ and for all $c \in W$, if $h_1 \perp_c h_2$, then $c \in C$.*

Proof. "$\Rightarrow$": Assume that $\bigcup\lambda$ is blatantly inconsistent, and let a witnessing pair be $\tau_1 = e \rightarrowtail H_1 \in T_1 \in \lambda$ and $\tau_2 = e \rightarrowtail H_2 \in T_2 \in \lambda$, $H_1 \neq H_2$. Let

$h_1 \in H(T_1)$ and $h_2 \in H(T_2)$; then $e \in h_1 \cap h_2$ and $\Pi_e\langle h_1\rangle = H_1 \neq H_2 = \Pi_e\langle h_2\rangle$, i.e., $h_1 \perp_e h_2$.

"$\Leftarrow$": Assume that there are $h_1, h_2 \in H(\lambda)$ for which $h_1 \perp_e h_2$ for some $e \in E$. Then there must be $T_1, T_2 \in \lambda$ for which $(e \rightarrowtail \Pi_e\langle h_i\rangle) \in T_i$ $(i = 1, 2)$, and $\Pi_e\langle h_1\rangle \neq \Pi_e\langle h_2\rangle$. As $\bigcup\lambda \supseteq T_1 \cup T_2$, the set $\bigcup\lambda$ contains two different transitions with initial e and is, therefore, blatantly inconsistent. □

This fact says that if a non-contextual instruction set λ is truly counterfactual (that is, if it has at least two different elements), then the difference is due to what happens outside of E: the corresponding histories split at elements of C only. As we will see, contextual instruction sets relax this condition by allowing that the corresponding histories may split at members of C *and* at members of E.

In our discussion of generic instruction sets we remarked that such instruction sets can be built from any consistent set of transitions. A similar, but more specific result also holds for non-contextual instruction sets; it is worth spelling out in detail.

Fact 8.9. *Let $T \subseteq \tilde{T}$ be a consistent set of transitions with initials $E_T \subseteq E$. Then there is some non-contextual instruction set $\lambda \in \mathfrak{I}_n$ for which $E_T \subseteq E_\lambda$ and $\lambda(E_T) = T$.*

Proof. The given T can be extended to a maximal consistent set $T^* \in S$. Note that $\lambda_0 =_{\text{df}} \{T^*\}$ fulfills the conditions (1) and (2) of Def. 8.11: (1) holds by construction, and (2) is trivial as T^* is consistent. There is a maximal extension of λ_0, $\lambda \in \mathfrak{I}_n$, which retains conditions (1) and (2). As a superset of λ_0, for λ we have $E_\lambda \supseteq E_{\lambda_0} \supseteq E_T$. And by the choice of λ_0 and by (1), it must be that $\lambda(E_T) = \lambda_0(E_T) = T$. □

Here is another simple fact about non-contextual instruction sets.

Fact 8.10. *Let $\lambda \in \mathfrak{I}_n$. (1) The set of transitions $\bigcup\lambda = \lambda(E_\lambda)$ is downward closed, i.e., if $\tau \in \bigcup\lambda$ and for some $\tau' \in \tilde{T}$ we have $\tau' \prec \tau$, then $\tau' \in \bigcup\lambda$. (2) For $T \in S$, we have $T \in \lambda$ iff $T \subseteq \bigcup\lambda$.*

Proof. (1) This claim follows directly from Def. 8.11, as any $T \in S$ is downward closed by maximality.

(2) "$\Rightarrow$": If $T \in \lambda$, then clearly any element of T is in $\bigcup\lambda$.

"$\Leftarrow$": Let $T \subseteq \bigcup\lambda$, so $\bigcup(\lambda \cup \{T\}) = \bigcup\lambda$. As $\bigcup\lambda$ fulfills the conditions (1) and (2) of Def. 8.11, $\bigcup(\lambda \cup \{T\})$ fulfills these conditions as well, and so $T \in \lambda$ follows by maximality of λ. □

The definition of a non-contextually extended structure follows our template of structure extensions from Section 8.3.1. The definitions of matching and of lifted histories (Defs. 8.5 and 8.6) remain unaltered. We write out the definition of the extended structure for the sake of completeness.

Definition 8.12. Let $\langle W, <, e^*, E, C\rangle$ be a BST_{92} surface structure and let $\mathfrak{I}_n$ be its set of non-contextual instruction sets. The corresponding *non-contextual extension* is the extended structure $\langle W', <'\rangle$ of Def. 8.7, replacing $\mathfrak{I}_g$ by $\mathfrak{I}_n$.

Theorem 8.1 applies and shows that such $\langle W', <'\rangle$ is again a BST_{92} structure. Furthermore, the histories $h \in \mathrm{Hist}(W)$ and $h' \in \mathrm{Hist}(W')$ are related by Lemma 8.1. For the splitting of these histories, clauses (1)–(5) of Lemma 8.2 are relevant; clause (6) is excluded as for $\lambda \in \mathfrak{I}_n$, the set of transitions $\bigcup\lambda$ cannot be blatantly inconsistent by Def. 8.11. This implies that in an extended structure based on non-contextual instruction sets, there can be no MFB between outcomes of members of E any more, independent of such MFB in the surface structure: the respective copies of members of E are no longer choice points. There is, however, a new choice point $\langle e^*, \emptyset\rangle \in W'$, and the crucial question is whether an extended structure will exhibit C/Ext independence. As we will show in our discussion of the GHZ experiment in Section 8.3.3, C/Ext independence can fail even if the surface structure is C/E independent. So, while the extension by non-contextual instruction sets is well-defined for any BST_{92} surface structure, it will not always be satisfactory (see Theorem 8.3).

8.3.2.2 Contextual instruction sets

Contextual instruction sets relax a constraint on the non-contextual ones: the prescribed outcome for some $e \in E$ need not be unique, but may depend on the context provided by outcomes of C. In contrast to generic instruction sets, for which there are no constraints, contextual instruction sets *are*, however, constrained: if they specify different outcomes for some $e \in E$, there must also be a difference in the outcome of some $c \in C$. In the following definition, this is spelled out as the condition of C-splitting.

Definition 8.13 (Contextual instruction sets)**.** Let $\langle W,<,e^*,E,C\rangle$ be a BST_{92} surface structure. We say that $\lambda \subseteq S$ is a *contextual instruction set for* $\langle W,<,e^*,E,C\rangle$ iff λ is maximal with respect to the following condition of *C-splitting*: For any two $T_1,T_2 \in \lambda$, $T_1 \neq T_2$, the T_1- and T_2-histories split at a member of C, i.e.,

$$\forall T_1,T_2 \in \lambda\ [T_1 \neq T_2 \to \forall h_1 \in H(T_1)\,\forall h_2 \in H(T_2)\,\exists c \in C\ [h_1 \perp_c h_2]].$$

We write $\mathfrak{I}_c$ for the set of all contextual instruction sets. We call an instruction set $\lambda \in \mathfrak{I}_c$ *properly contextual* iff it provides inconsistent instructions for at least one $e \in E$, i.e., iff $\bigcup \lambda$ is blatantly inconsistent.

Contextual instruction sets can be constructed starting from any consistent set of transitions. That is, we have the following analogue of Fact 8.9:

Fact 8.11. *Let $T \subseteq \tilde{T}$ be a consistent set of transitions. There is some contextual instruction set $\lambda \in \mathfrak{I}_c$ for which $T \subseteq \bigcup \lambda$.*

Proof. Left as Exercise 8.3. □

It is also clear that contextual instruction sets are downward closed (see Fact 8.10(1)).

Many further remarks on contextual instruction sets parallel our remarks for non-contextual ones. The definition of a contextually extended structure follows our template of structure extensions from Section 8.3.1. The definitions of matching and of lifted histories (Defs. 8.5 and 8.6) remain unaltered. We write out the definition of the extended structure for the sake of completeness.

Definition 8.14. Let $\langle W,<,e^*,E,C\rangle$ be a BST_{92} surface structure and let $\mathfrak{I}_c$ be its set of contextual instruction sets. The corresponding *contextual extension* is the extended structure $\langle W',<'\rangle$ of Def. 8.7, replacing $\mathfrak{I}_g$ by $\mathfrak{I}_c$.

Theorem 8.1 applies and shows that such $\langle W',<'\rangle$ is again a BST_{92} structure. Furthermore, the histories $h \in \text{Hist}(W)$ and $h' \in \text{Hist}(W')$ are related by Lemma 8.1. For the splitting of these histories, all of the clauses (1)–(6) of Lemma 8.2 may be relevant. If there are properly contextual $\lambda \in \mathfrak{I}_c$, then in an extended structure there are choice points, apart from those based on elements of C, of the form $\langle e,\lambda\rangle \in W'$ (besides $\langle e^*,\emptyset\rangle \in W'$), so that the crucial question of *C/Ext* independence becomes more complex. As we will show in our discussion of the GHZ experiment in Section 8.3.3, it can happen that a surface structure is C/E independent, while the extended structure violates C/Ext independence. So, while the extension by contextual instruction sets is well-defined for any BST_{92} surface structure, it will not always be satisfactory (see Theorem 8.4).

8.3.2.3 On the interrelation of different types of instruction sets

We have provided four different definitions of instruction sets, each singling out a unique set of subsets of S for a given BST_{92} surface structure $\langle W, <, e^*, E, C\rangle$, with the proviso that superdeterministic instruction sets are defined only if $C = \emptyset$. In that case, it turns out that the non-contextual and the contextual instruction sets are singletons of elements of S, thus coinciding with the superdeterministic instruction sets.

Fact 8.12. *Let $\mathscr{W} = \langle W, <, e^*, E, C\rangle$ be a BST_{92} surface structure with $C = \emptyset$, so that E is the set of* all *choice points in $\mathscr{W}$. Then the superdeterministic, non-contextual, and contextual instruction sets coincide, i.e., $\mathfrak{I}_s = \mathfrak{I}_n = \mathfrak{I}_c$.*

Proof. There can be no contextual instruction set $\lambda \in \mathfrak{I}_c$ that contains more than one member of S, as the condition of external splitting is impossible to fulfill given $C = \emptyset$. Thus, the $\lambda \in \mathfrak{I}_c$ are singletons of elements of S. By Fact 8.8, the same holds for non-contextual instruction sets $\lambda \in \mathfrak{I}_n$. So, any instruction set, superdeterministic, non-contextual, or contextual, is a singleton of an element $T \in S$. □

The corresponding superdeterministic extended structure has already been discussed in Section 8.3.1.1: by Lemma 8.2, it contains just one choice point, $\langle e^*, \emptyset\rangle$, which has as many outcomes as there are histories in the surface structure. Thus, in this superdeterministic extended structure, everything is decided at $\langle e^*, \emptyset\rangle$. By the assumption $C = \emptyset$, there is no agent-induced indeterminism in the surface structure, nor is there such indeterminism in the extended structure.

In the general case, the four notions of instruction sets single out different subsets of S, but they nest in the sense that instruction sets of a more restricted type can be extended to instruction sets of a looser type. The following lemma spells this out in formal detail.

Lemma 8.3 (Nesting of types of instruction sets)**.** *The sets of instruction sets $\mathfrak{I}_s$, $\mathfrak{I}_n$, $\mathfrak{I}_c$, and $\mathfrak{I}_g$ nest in the following way: (1) If $C = \emptyset$, so that $\mathfrak{I}_s$ is defined,* *we have $\mathfrak{I}_s = \mathfrak{I}_n = \mathfrak{I}_c \subseteq \mathfrak{I}_g$. (2) For each $\lambda \in \mathfrak{I}_n$, there is some $\lambda' \in \mathfrak{I}_c$ for which $\lambda \subseteq \lambda'$. (3) For each $\lambda \in \mathfrak{I}_c$, there is some $\lambda' \in \mathfrak{I}_g$ for which $\lambda \subseteq \lambda'$.*

Proof. (1) The two equalities have been established via Fact 8.12. The inclusion is trivial, as $\mathfrak{I}_g$ is the set of *all* subsets of S.

(2) It is easy to prove that each $\lambda \in \mathfrak{I}_n$ fulfills the condition of C-splitting: As $\bigcup\lambda$ contains no blatantly inconsistent transitions, there can be no T_1, T_2 such that histories $h_1 \in H(T_1)$, $h_2 \in H(T_2)$ split at a member of E at all. (See also Fact 8.8.)

(3) As any $\lambda \in \mathfrak{I}_c$ is a subset of S, λ itself is already a member of $\mathfrak{I}_g$, which comprises *all* subsets of S. □

8.3.3 Instruction sets for GHZ

In this section we put the concept of structure extension to work on a well-known example. We will use it to analyze the three-particle GHZ experiment (Greenberger et al., 1989) in the form presented by Mermin (1990). In this experiment, a source emits triples of particles that fly to remote measurement stations. Each triple is in the quantum state

$$\lambda = 1/\sqrt{2}(|+\rangle\,|+\rangle\,|+\rangle - |-\rangle\,|-\rangle\,|-\rangle), \tag{8.1}$$

which is a vector in the 8-dimensional Hilbert space $\mathscr{H} = \mathscr{H}_2 \otimes \mathscr{H}_2 \otimes \mathscr{H}_2$. At each station i ($i = 1,2,3$), the experimenter can choose to measure the spin projection on the direction x_i or y_i, where x_i is the direction perpendicular to the plane of flight, whereas y_i is in the plane of flight, perpendicular to the trajectory of the ith particle. For each measurement setting, x_i or y_i, there are two alternative possible outcomes, x_i^+ and x_i^-, or y_i^+ and y_i^-.

Taking the description of the experimenter's choices and alternative measurement outcomes at face value, there are two layers of choice points: the experimenter's choices at the three stations correspond to $C = \{c_1, c_2, c_3\}$, and the measurements at the three stations, with the two possible directions x_i and y_i, correspond to $E = \{x_1, y_1, x_2, y_2, x_3, y_3\}$. The choice of direction is obviously prior to the respective measurements, i.e., $c_i < x_i$ and $c_i < y_i$ ($i = 1,2,3$). We can select a location for e^* that is in the past of all of E, yet *SLR* to all of C, which corresponds to the spatio-temporal arrangement of the particle source. Thus we have specified elements e^*, E, and C for a valid BST_{92} surface structure.

In our BST_{92} analysis we represent the indeterminism in the experiment by transitions, first from an experimenter's indecision c_i at station i to one of the selected settings, x_i, or y_i, and then from the measurement with a given setting, say x_i, to one of its outcomes, x_i^+ or x_i^-. To simplify the notation, we will write c_1^x for the transition $c_1 \rightarrowtail \Pi_{c_1}\langle x_1\rangle$ and y_2^- for $y_2 \rightarrowtail \Pi_{y_2}\langle y_2^-\rangle$, and so on.[16] So we have the following indeterministic transitions in our structure:

$$c_1^x, c_1^y,\ c_2^x, c_2^y,\ c_3^x, c_3^y \quad \text{(transitions due to agents);}$$
$$x_1^+, x_1^-, y_1^+, y_1^-,\ x_2^+, x_2^-, y_2^+, y_2^-,\ x_3^+, x_3^-, y_3^+, y_3^- \quad \text{(transitions due to Nature).}$$

Figure 8.2 provides a schematic picture of the two layers of indeterminism.

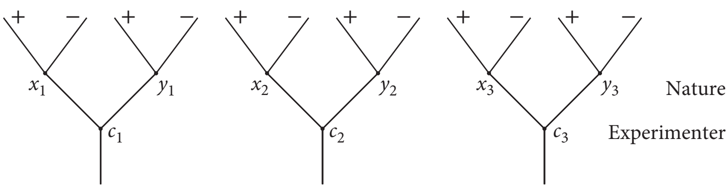

Figure 8.2 Schematic diagram of the GHZ experiment. Choice points c_i are the experimenters' selections of measurement directions x_i or y_i, and x_i (or y_i) are choice points with possible outcomes '+' or '−' determined by Nature ($i = 1, 2, 3$).

To exclude the possibility of physically transmitted causal influences in this experiment, the following *SLR* relations are assumed: (1) between the different experimenter's choices (e.g., $c_1\,SLR\,c_2$), (2) between an experimenter's choice and a remote measurement (e.g., $c_1\,SLR\,x_2$), and (3) between measurements in different stations (e.g., $x_1\,SLR\,y_3$).

So far, there is nothing strange or curious about this experiment. You could think of the source as emitting triples of marbles in the three directions, with the experimenters selecting to measure their size or color, for example. The quantum-mechanical strangeness of the GHZ experiment lies in the fact that only some, but not all triples of *joint* transitions to experimental outcomes such as $\langle x_1^-, y_2^-, y_3^+\rangle$ are possible, even though all *individual* transitions to the outcomes (such as x_2^-) are possible. Given the

[16] That is, in the latter case we use the expression 'y_2^-' both for the transition from initial y_2 to its '−' outcome and for the '−' outcome itself. It will always be clear from context what is meant.

state of Eq. 8.1 and directions x_i, y_i, quantum mechanics assigns non-zero probabilities to 48 out of the 64 possible triples. But for the 16 other triples, quantum mechanics assigns probability zero. In our current modal framework we interpret these probabilistic verdicts as verdicts about possibility and impossibility: triples of the first kind are consistent, while triples of the second kind are inconsistent.[17] Read in this way, quantum mechanics implies that there are modal correlations among the measurement outcomes. These can be described by the following two simple rules:

(yy) A triple with two *y*'s and one *x* is consistent iff it includes an even number of minuses.
(xxx) A triple with all three *x*'s is consistent iff it includes an odd number of minuses.

The rules imply, among other things, that there is no history that includes the transitions x_1^+, x_2^+, and x_3^+ together. This lacking history indicates combinatorial funny business, as the set of transitions $T =_{\text{df}} \{x_1^+, x_2^+, x_3^+\}$ is combinatorially consistent (for the initials x_1, x_2, x_3, we have $x_i \, SLR \, x_j$ for each pair i, j, $i \neq j$), but $H(T) = \emptyset$. As we indicated, and as one can check via the rules (yy) and (xxx), MFB in this structure is in fact ubiquitous. In total, 64 combinations of transitions are combinatorially consistent (2^3 combinations of initials by 2^3 outcome combinations), but only 48 of these satisfy both rules, so there are 16 cases of MFB.

Via the above discussion we have defined our BST_{92} surface structure for GHZ:[18]

Definition 8.15. The *surface structure for GHZ* is the BST_{92} surface structure $\langle W, <, e^*, E, C\rangle$ discussed above, in which $C = \{c_1, c_2, c_3\}$, $E = \{x_1, y_1, x_2, y_2, x_3, y_3\}$, and $e^* < E$ and $e^* \, SLR \, C$. There are the further *SLR* relations noted above (e.g., $c_1 \, SLR \, c_2$, $c_1 \, SLR \, y_2$, $x_1 \, SLR \, y_2$). The structure has exactly those 48 histories that satisfy the two given rules (yy) and (xxx).

We can now study structure extensions.

[17] The experiment can also be analyzed probabilistically in BST—see Section 8.4.3 for a BST approach to probabilistic hidden variable models.
[18] In order not to bury the important message of our discussion under a mound of formalism, we do not define the ordering in the structure fully formally, relying instead on the specification of the relevant *SLR* relations. It is possible to provide a fully formal specification, e.g., as a Minkowskian Branching Structure as discussed in Chapter 9.1.

8.3.3.1 The superdeterministic extension

We mention, only to leave it to the side, the issue of a superdeterministic extension. As $C \neq \emptyset$, our given BST surface structure does not allow for a superdeterministic extension according to Def. 8.8, but it could be modified in the following way: move e^* to the past such that the new $e'^* < C$, and then set $E' =_{\text{df}} E \cup C$ and $C' =_{\text{df}} \emptyset$. The structure $\langle W, <, e'^*, E', C' \rangle$ is a BST_{92} surface structure that allows for a superdeterministic extension. In that extension, the 48 histories then split at $\langle e'^*, \emptyset \rangle$, and any outcome of $\langle e'^*, \emptyset \rangle$ fixes both the experimenter's choice of settings (x_i or y_i) and the respective measurement outcomes (x_i^+ or x_i^-, for example). This extended structure is philosophically and physically unilluminating: it summarizes which individual runs of the experiment are possible, but it does not describe the experiment as an experiment any more. A useful extension needs to stick with the given surface structure in which the experimenter's choices of measurement directions are formally distinguished from Nature's choices of the measurement outcomes.

8.3.3.2 C/E independence

Before we move to the construction of non-contextual and contextual structure extensions for GHZ, we have to discuss the crucial issue of C/E independence. When specifying the rules (yy) and (xxx), we noted that these rules amount to the introduction of MFB among the measurement outcomes—there are subsets of $\tilde{T}$, such as $\{x_1^+, x_2^+, x_3^+\}$, that are combinatorially consistent but inconsistent. What about MFB involving choice points for the selection of settings c_i and for the selection of measurement outcomes x_j? By the given SLR relations, and as the choice points x_i and y_i are incompatible ($i = 1, 2, 3$), a combinatorially consistent set that includes transitions both from C and from E has to have one of the following schematic forms (where α, β, γ stand for x or y, and m, n stand for $+$ or $-$), leading to a total of 144 sets:

$$c_1^\alpha, c_2^\beta, \gamma_3^m;\ \ c_1^\alpha, \beta_2^m, c_3^\gamma;\ \ \alpha_1^m, c_2^\beta, c_3^\gamma;\ \ \alpha_1^m, \beta_2^n, c_3^\gamma;\ \ \alpha_1^m, c_2^\beta, \gamma_3^n;\ \ c_1^\alpha, \beta_2^m, \gamma_3^n.$$

That is, such a set either specifies two outcomes for the experimenter's selection of parameters at two of the stations (e.g., c_1^x, c_2^x) and one outcome of a specific measurement at the third station (e.g., x_3^+), or it specifies one outcome of parameter selection (e.g., c_3^x) and two outcomes of specific measurements at the other two stations (e.g., x_1^+, x_2^+). Obviously, by the downward closure of histories, if such a set includes a specific outcome such as x_3^+, this implies the corresponding choice of measurement setting, in this case, c_3^x. By Lemma 5.2, however, such a downward extension preserves both consistency and combinatorially consistency, so we can save the labor of writing out the additional transitions (e.g., we stick to our c_1^x, c_2^x, x_3^+ in place of the longer $c_1^x, c_2^x, c_3^x, x_3^+$).

Given the structure of the combinatorially consistent sets mixing transitions from C and from E and the structure of the rules (yy) and (xxx), it is easy to see that these mixed sets are all consistent, so that there is no case of MFB among these sets. That is, the GHZ surface model satisfies C/E independence:

Fact 8.13. *The GHZ surface structure defined in Def. 8.15 satisfies the condition of C/E independence of Def. 8.3.*

Proof. We check three exemplary cases. (1) For $T_1 = \{c_1^x, c_2^x, x_3^+\}$, we need to take heed of the (xxx) rule. The set of outcomes x_1^-, x_2^+, x_3^+ satisfies that rule, so there is a corresponding history $h^{x-x+x+} \in H(T_1)$. (2) For $T_2 = \{c_1^x, y_2^-, c_3^y\}$, we need to take heed of the (yy) rule. The set of outcomes x_1^-, y_2^-, y_3^+ satisfies that rule, so there is a corresponding history $h^{x-y-y+} \in H(T_2)$. (3) For $T_3 = \{x_1^+, x_2^-, c_3^x\}$, we again need to take heed of the (xxx) rule. The set of outcomes x_1^+, x_2^-, x_3^+ satisfies that rule, so there is a corresponding history $h^{x+x-x+} \in H(T_3)$. The structural point is that a set that includes transitions both from C and from E always leaves at least one outcome $\pm$ unspecified, and so if one of the rules (yy) or (xxx) applies, that outcome can be chosen accordingly. The upshot is that any combinatorially consistent set of transitions involving both C and E as initials is consistent. □

This fact may raise the hope that an extension of the GHZ surface structure that removes the troublesome cases of MFB could still satisfy the companion notion of C/Ext independence (Def. 8.10). The crucial question is whether this is possible and, if so, for which type of instruction sets.

8.3.3.3 Non-contextual instruction sets for GHZ

The idea of an instruction set is to provide guidance for Nature's choice points $e \in E$ in a way that is independent of the experimenters' choices at $c \in C$. Non-contextual instruction sets are such that an instruction set $\lambda \in \mathfrak{I}_n$ specifies exactly one outcome for each initial in its domain, so that λ can also be read as a partial function from E to $\tilde{T}$. The set-up of GHZ is simple in this respect, because it is layered: the experimenter selects three parameters x_i or y_i ($i = 1,2,3$), and then Nature provides one of the possible joint outcomes. The MFB is among the members of $\tilde{T}$ only, as we have just shown (Fact 8.13).

Intuitively, one would expect a non-contextual instruction set to be a function that specifies maximal information, meaning that it specifies one outcome each for each of $x_1, \ldots, y_3$. The sobering fact is that this is not possible in a way that satisfies the GHZ rules.

Fact 8.14. *There is no way to specify one outcome each for the six possible measurements $x_1, \ldots, y_3$ in such a way that both rules (yy) and (xxx) are fulfilled.*

Proof. There are eight six-element sets of outcomes that satisfy the (yy) rule: in total, there are $2^6 = 64$ possible combinations, and the (yy) rule cuts this number down by a factor of $2^3 = 8$, as it imposes three parity constraints. We list the eight satisfactory combinations here. The first line gives those for which the outcomes of x_1 and y_1 agree; the second line gives those for which x_1 and y_1 disagree.

$$\{x_1^+ x_2^+ x_3^+ y_1^+ y_2^+ y_3^+\}\ \{x_1^+ x_2^- x_3^- y_1^+ y_2^- y_3^-\}\ \{x_1^- x_2^+ x_3^- y_1^- y_2^+ y_3^-\}\ \{x_1^- x_2^- x_3^+ y_1^- y_2^- y_3^+\}$$
$$\{x_1^+ x_2^- x_3^- y_1^- y_2^+ y_3^+\}\ \{x_1^+ x_2^+ x_3^+ y_1^- y_2^- y_3^-\}\ \{x_1^- x_2^- x_3^+ y_1^+ y_2^+ y_3^-\}\ \{x_1^- x_2^+ x_3^- y_1^+ y_2^- y_3^+\}.$$

As one can see by inspection, none of these eight sets satisfies the (xxx) rule, as the number of minuses on the x-outcomes is always zero or two. $\square$

Before this background, it is interesting to ask what the non-contextual instruction sets look like for GHZ and which effects it has that these sets are not maximally specific. Our result is that a non-contextual extension for GHZ introduces C/Ext dependence despite the surface structure's C/E independence.

Instruction sets $\lambda \in \mathfrak{I}_n$ are subsets of S maximal with respect to two conditions spelled out in Def. 8.1: (1) such a set must provide a unanimous instruction for each $e \in E$ covered, so that $\bigcup \lambda$ is not blatantly inconsistent, and (2) such a set must actually provide instructions for consistent sets of initials (i.e., if $E_0 \subseteq E_\lambda$ is consistent, then so is the corresponding set of transitions $\lambda(E_0)$).

The set S has 48 members, each corresponding to a history in the surface structure (e.g., $T_0 =_{\mathrm{df}} \{x_1^+, x_2^+, x_3^-\} \in S$). It is instructive to see, for example, what happens if one extends T_0 to a non-contextual instruction set λ. We provide one illustration; by Fact 8.14, all alternative attempts at constructing an instruction set will end up in a similar predicament. Consider first, as a warm-up, a construction that is constrained only by condition (1). Basically, the task is to select outcomes for y_1, y_2, and y_3, as the outcomes for the x's are already fixed. So let us try, for example, y_1^+, y_2^-, and y_3^+. The members of S that correspond to this selection are the following:

$$\{x_1^+, x_2^+, x_3^-\}, \{x_1^+, x_2^+, y_3^+\}, \{x_1^+, y_2^-, x_3^-\}, \{y_1^+, x_2^+, x_3^-\},$$
$$\{y_1^+, x_2^+, y_3^+\}, \{y_1^+, y_2^-, x_3^-\}, \{y_1^+, y_2^-, y_3^+\}.$$

What is missing from this list is a set that provides guidance as to the selection of settings x_1, y_2, y_3—the set that would have to be included, as forced by no blatant inconsistency, is $\{x_1^+, y_2^-, y_3^+\}$, but this set violates the (yy) rule, and so there is no history in the GHZ surface structure that includes it.

For the real instruction sets, we have to consider both conditions (1) and (2). Obviously, the resulting λ coming from T_0 is a subset of the set displayed above. The result is

$$\lambda = \{\{x_1^+, x_2^+, x_3^-\}, \{x_1^+, x_2^+, y_3^+\}, \{y_1^+, x_2^+, x_3^-\}, \{y_1^+, x_2^+, y_3^+\}\}. \tag{8.2}$$

For example, the set $\{x_1^+, y_2^-, x_3^-\}$, which is the third one displayed above, cannot be included because given the two previous sets $\{x_1^+, x_2^+, x_3^-\}$ and $\{x_1^+, x_2^+, y_3^+\}$, adding it violates condition (2): the set of initials $\{x_1, y_2, y_3\}$ would otherwise be a consistent subset of the set of initials of these three sets of transitions, and the corresponding set of transitions, $\{x_1^+, y_2^-, y_3^+\}$, violates the (yy) rule. The reader is invited to check that the given λ cannot be extended while still satisfying both conditions (1) and (2); see Exercise 8.5.

Note that when viewed as a partial function, λ is defined on five out of the six members of E, viz., x_1, x_2, x_3, y_1, y_3. By Fact 8.14, that is the maximum domain for a consistent partial function. Lacking an outcome instruction for y_2, our displayed instruction set λ does not provide instructions for four out of the eight possible selections of settings. Thus, the extended structure violates C/Ext independence: if Nature selects our instruction set λ at $\langle e^*, \emptyset \rangle$, then the experimenter must be prohibited, among other things, from choosing the parameters y_1, y_2, and y_3. At any rate, the choice of y_2

must be prevented somehow. Formally, to show how our definitions apply, we can spell out this result in the form of the following "no go" Theorem:

Theorem 8.3. *There are no non-contextual hidden variables for the GHZ experiment.*

More precisely: For the GHZ surface structure $\langle W, <, e^*, E, C\rangle$ *defined in Def. 8.15, which is* C/E*-independent, the extension by non-contextual instruction sets* $\mathfrak{I}_n$ *results in an extended structure* $\langle W', <'\rangle$ *that violates* C/Ext *independence (Def. 8.10).*

Proof. The C/E-independence of the surface structure has been shown via Fact 8.13. To prove C/Ext dependence formally, consider the non-contextual instruction set of Eq. (8.2) discussed above,

$$\lambda = \{\{x_1^+, x_2^+, x_3^-\}, \{x_1^+, x_2^+, y_3^+\}, \{y_1^+, x_2^+, x_3^-\}, \{y_1^+, x_2^+, y_3^+\}\}.$$

The set of transitions in W' (where we abbreviate the outcomes in a mnemonic way)

$$T = \{\langle e^*, \emptyset\rangle \rightarrowtail H^\lambda, \langle c_1, \emptyset\rangle \rightarrowtail H^y, \langle c_2, \emptyset\rangle \rightarrowtail H^y, \langle c_3, \emptyset\rangle \rightarrowtail H^x\}$$

is combinatorially consistent, as all initials are pairwise SLR, but it is inconsistent, as no history h matching λ includes the transitions c_1^y, c_2^y, and c_3^x. □

That is, the extended structure is conspiratorial, and the process of extending the surface structure turns out to be a cure worse than the disease: from an initially hard-to-understand case of modal funny business without any Nature-experimenter conspiracies, the process of structure extension has created a structure in which the modal funny business between Nature's outcomes is removed at the expense of introducing modal funny business between Nature's single choice and the experimenter's choice of measurement parameters. This is certainly not a satisfactory kind of structure extension.

8.3.3.4 Contextual instruction sets for GHZ

The failure of the non-contextual approach to GHZ motivates the use of contextual instruction sets. These are subsets of S that are less strictly constrained than non-contextual sets, as the instructions given for Nature's choice points may depend on the measurement context. For the GHZ set-up, such a measurement context consists of exactly one choice of direction x or y at each station $i = 1, 2, 3$. As with the previous case of non-contextual instruction sets, it is instructive to check in formal detail what the structure extension turns out to be like.

By Def. 8.13, a contextual instruction set is a subset of S that is maximal with respect to what is there called C-splitting. It is easy to check that the sets of transitions $T_1, T_2 \in S$ are C-splitting exactly if they differ in at least one measurement parameter. Thus, for example, the sets $\{x_1^+, x_2^+, x_3^-\}$ and $\{x_1^+, x_2^+, y_3^+\}$ are C-splitting, but the sets $\{x_1^+, x_2^+, x_3^-\}$ and $\{x_1^-, x_2^+, x_3^+\}$ are not. Any selection of parameters allows for a consistent choice by Nature, as one can see by inspecting the (yy) and (xxx) rules, and there are no further constraints on contextual instruction sets. Thus, a contextual instruction set for GHZ has exactly eight elements, one for each possible combination of parameter choices, $x_1x_2x_3, x_1x_2y_3, \ldots, y_1y_2y_3$. A fortiori, such an instruction set provides information pertaining to all six possible measurement choice points $E = \{x_1, \ldots, y_3\}$. By Fact 8.14, it is impossible to specify one outcome each for all six possible measurements in such a way that the (yy) and (xxx) rules are satisfied. Therefore, any $\lambda \in \mathfrak{I}_c$ has to be properly contextual, i.e., $\bigcup \lambda$ must be blatantly inconsistent. Here is one exemplary $\lambda \in \mathfrak{I}_c$:

$$\begin{aligned}\lambda = \{&\{x_1^-, x_2^-, x_3^-\}, \{x_1^-, x_2^-, y_3^-\}, \{x_1^-, y_2^-, x_3^-\}, \{x_1^- y_2^- y_3^+\},\\ &\{y_1^+, x_2^-, x_3^-\}, \{y_1^+, x_2^-, y_3^-\}, \{y_1^+, y_2^-, x_3^-\}, \{y_1^-, y_2^-, y_3^+\}\}.\end{aligned} \tag{8.3}$$

To see that λ is properly contextual, note that

$$\bigcup \lambda = \{x_1^-, x_2^-, x_3^-, y_1^-, y_2^-, y_3^+, y_3^-\},$$

so that there are two different outcomes prescribed for $y_3 \in E$, depending on the context (compare the second and the fourth elements displayed, $\{x_1^-, x_2^-, y_3^-\}$ and $\{x_1^-, y_2^-, y_3^+\}$). By Lemma 8.2(6), this means that the element $\langle y_3, \lambda\rangle \in W'$ in the contextual extension $\langle W', <'\rangle$ remains a choice point.

Now given our λ and looking at the instance of blatant inconsistency just described, one sees that the outcome of y_3 depends on which measurement parameter is chosen at station 2. (Our λ also includes a case of sensitivity for the choice of station 1.) Such a dependence looks as if the structure extension has introduced novel cases of MFB between experimenters' choices and some of Nature's choices, thus showing that the attempt to provide instruction sets for GHZ in the contextual way is also futile. The following "no go" Theorem spells out in formal detail that this is indeed the case.

Theorem 8.4. *There are no contextual hidden variables for the GHZ experiment.*

More precisely: For the GHZ surface structure $\langle W, <, e^*, E, C\rangle$ defined in Def. 8.15, which is C/E-independent, the extension by contextual instruction sets $\mathfrak{I}_c$ results in an extended structure $\langle W', <'\rangle$ that violates C/Ext independence (Def. 8.10).

Proof. The C/E-independence of the surface structure has been shown via Fact 8.13. To prove C/Ext dependence formally, consider the contextual instruction set λ of Eq. (8.3) discussed above and the following set of transitions in W' (where we abbreviate the outcomes in an obvious way):

$$T = \{\ \langle c_1, \emptyset\rangle \rightarrowtail H^x, \langle c_2, \emptyset\rangle \rightarrowtail H^x, \langle y_3, \lambda\rangle \rightarrowtail H^+\}.$$

That set is combinatorially consistent, as all initials are pairwise SLR, but it is inconsistent: a history h' including these transitions would have to include $\langle y_3, \lambda\rangle$, so by Fact 8.3, there would have to be some $h \in \text{Hist}(W)$ matching λ for which $h' = \varphi_\lambda(h)$. Furthermore, $\{x_1, x_2, y_3\} \subseteq h$, by the members of T. Now the only history in W that matches λ and which contains these three measurement initials is h^{x-x-y-}, for which the initial y_3 has outcome '$-$'; but by the third element of T, y_3 would have to have outcome '$+$'. □

That is, similarly to the case of non-contextual instruction sets, the structure extended by contextual instruction sets is conspiratorial, and the process of extending the surface structure turns out to be a cure worse than the disease: from an initially hard-to-understand case of modal funny business without any Nature-experimenter conspiracies, the process of structure extension has created a structure in which the modal funny business between Nature's outcomes is removed at the expense of introducing modal funny business between one of Nature's remaining choice points and the experimenter's choice of measurement parameters. Again, this is certainly not a satisfactory kind of structure extension.

8.3.4 Summary of the BST approach to modal structure extensions

In this part of our investigation of quantum correlations and their modeling, we chose an approach that is specific to BST, as it uses the formally well-defined notion of modal funny business to motivate structure extensions, to define them, and to gauge their success.

We provided a template for the general case of structure extensions, showing how these create BST_{92} structures from BST_{92} surface structures. Our framework allowed us to specify exactly what a superdeterministic (or conspiratorial) extension is, and it allowed us to discuss the issue of dependence and independence between experimenters' and Nature's choices in formal detail. We spelled out general conditions of adequacy for structure extensions, which require that such an extension should not introduce novel cases of modal funny business involving choices of the experimenters and outcomes provided by Nature: the objective is to remove surface MFB by (1) modifying choices given by Nature, while (2) leaving experimenters' choices intact.

We defined the well-known approaches of non-contextual and contextual structure extensions in the BST framework. In order to put our definitions to the test, we tackled the GHZ experiment, which is known not to admit sensible hidden variable extensions. We could reproduce these findings in the form of two formally well-specified no go-theorems.

The upshot is that BST allows for structure extensions, both contextual and non-contextual, that are always well-defined. Whether they are satisfactory or not, however, depends on the set-up in question. There are quantum set-ups that exhibit modal correlations that cannot be satisfactorily explained via structure extensions, neither non-contextual nor contextual ones.

8.4 Probabilistic correlations

Our task in the following sections is to investigate whether it is possible to explain *probabilistic* funny business (PFB) by invoking *probabilistic* hidden variables. This task sounds technical, but it concerns a non-technical philosophical problem. We will define the notion of extending BST_{92} structures with probabilistic hidden variables for the explanation of PFB. This notion is required to fulfill a number of intuitive desiderata concerning explanatoriness as well as the modal and spatio-temporal features of the phenomena in question. The technical task of defining structure extensions is motivated by the philosophical question of whether our world fulfils the mentioned desiderata. The literature on Bell-type theorems suggests an answer in the negative, and our analysis will support that verdict. Yet, the negative answer merely marks the point where our real investigation starts. We want to find out exactly which features of our world enforce this negative answer. Given BST, we have at our disposal a mathematically rigorous modal and spatio-temporal framework with gradable possibilities (propensities). Our hope is that this rich framework can deliver a precise answer as to why certain cases of PFB do not allow for explanatory extensions by probabilistic hidden variables. Before we embark on our analysis, we provide a brief discussion of the notion of probabilistic hidden variables.

8.4.1 Probabilistic hidden variables

In our analysis of modal correlations (Section 8.3), surface structures were extended by instruction sets, or deterministic hidden variables. Such instruction sets determine the measurement results, given the measurement settings. In contrast, probabilistic hidden variables work indeterministically: they do not fully determine a result, but instead they have a propensity to bring about a result. In the extreme, there can of course also be the "deterministically looking" propensities zero or one.

If the explanation of a case of PFB by means of probabilistic hidden variables is feasible, one can argue that the PFB is merely epistemic, as it will then be absent at the deepest level at which the hidden variables operate, and arise only at a shallower surface level, at which the distribution of hidden variables is averaged over. In studying this problem, we will use a format parallel to the one used in Section 8.3. We will assume that there is a surface structure, in this case a *probabilistic* BST_{92} surface structure harboring PFB, and we will ask whether this structure can be appropriately extended, so that the original PFB is removed in the resulting extended structure. As we appeal to probabilistic hidden variables, we do not aim to remove chanciness: instead, the aim is rather to remove the probabilistic surface correlations via an explanatory extension.

In contrast to our analysis in Section 8.3, which only made sense for structures in which MFB is present, in what follows we assume no MFB. This is in line with our analysis of probabilistic BST structures in Chapter 7. In particular, we assume NO MFB in the initial BST_{92} structure that represents phenomena with probabilistic funny business. When constructing an initial structure to represent phenomena that appear to lack some joint outcomes, like in the GHZ case discussed in Section 8.3.3, we will nevertheless represent such joint outcomes as possible in our initial BST_{92} structure, but we will assign them propensity zero to represent their absence. This move is very much in line with quantum mechanics, whose actual empirical predictions are always probabilistic.

The obvious aim of an explanatory probabilistic structure extension is to remove PFB. The extended structure, which must not harbor PFB, must of course adequately represent the probabilistic data encoded in the surface structure, including the probabilistic correlations. If an extension is successful, the surface probabilities will be recovered by collecting together the non-problematic probabilities in the extended structure. There are several constraints that limit the range of possible extensions. An important feature of our approach is the distinction between indeterminism resulting from experimental control and Nature's indeterminism, formally represented by sets of choice points C and E in the surface structure (see, e.g., Def. 8.2). Quantum mechanics predicts, and experiments confirm, probabilistic correlations between measurement outcomes, so that a structure extension would mainly operate on elements of E. However, this may also have some side-effects for the behavior of the choice points from C in the extended structure, and not all such side effects will be tolerable. An obvious desideratum is to require that the removal of PFB should not compromise the freedom of the experimenters. Another desideratum of this type is that in the extended structure, the experimental results should be independent from the remote choices of the experimenters—the motivation being that "remote" is identified with space-like relatedness and that faster-than-light signalling is prohibited. In the literature, the first desideratum is spelled out via a condition known as No Conspiracy, whereas the other desideratum goes by the name of Parameter Independence. (There is also a third condition that is frequently used in the literature, Outcome Independence: see Def. 8.22.) The independence that is invoked in these postulates can be interpreted modally or probabilistically, and thus the exact mathematical formulation is far from straightforward. The formal rigor of BST helps to state the conditions precisely and to analyze the proofs of theorems in which they are assumed. See Defs. 8.22 and 8.24 for our formulations.

We will analyze set-ups with differing degrees of complexity. The concrete way to analyze cases of PFB may depend on the complexity (or on our decision to acknowledge the complexity) of the set-up in question. The simplest case of PFB corresponds to an experiment with two fixed measurements performed in space-like separated regions. Such a set-up is represented in BST_{92} by a surface structure with a single case of PFB; that is, with a family of transitions with pairwise SLR initials (see our schematic representation of the EPR set-up, Figure 5.1, p. 107). As no choices of agents are present (there is no selection of alternative measurement settings), the set C of the experimenter-controlled choice points is empty, $C = \emptyset$. The set E, on the other hand, consists of the two SLR choice points corresponding to the initials of the transitions from the measurements to their possible results. It will not be surprising to learn that in such a case, a successful extension of the surface structure by probabilistic hidden variables is always possible.

More complex set-ups involve incompatible measurements, that is, measurements that cannot all occur together. Such set-ups can include the selection of the measurement settings by an experimenter at a choice point $c \in C$. The spatio-temporal ordering brings in yet another dimension of complexity, as there can be several experimenters, and their selections of parameters may be SLR. We may thus have several instances of PFB involving the results of pairwise incompatible experiments. BST represents such a set-up by several families of transitions, with the initials of the transitions from each family being SLR, and there always being some choice points responsible for the incompatibility of the initials of transitions from different families. These choice points can correspond to experimenter-induced selections of measurement settings. Similarly to our analysis in Section 8.3, we then face the task of spelling out, and of checking, the requirement of the independence of experimenters' and Nature's choices, thus making precise the general notion of C/E independence of Def. 8.1.

8.4.2 Extension of a probabilistic surface structure

Generally speaking, to explain probabilistic correlations between remote outcomes in BST, we first need to represent the phenomenon in question in a probabilistic BST surface structure, and then extend this structure in such a way that the extended structure harbors no correlations. In this section we discuss how an initial probabilistic BST_{92} structure, possibly containing probabilistic funny business, can be extended to a structure with a probabilistic hidden variable. As the journey to our target is rather lengthy, we offer a preview.

Starting with an initial BST_{92} structure with a case of PFB (but No MFB), our first task is order-theoretical: we introduce a multiplied structure that corresponds to the initial one, and we show it to be a BST_{92} structure. Apart from having some parts of the initial structure multiplied, this extended structure contains an extra choice point. The outcomes of this choice point together with propensities for these outcomes will serve as values of the hidden variable.[19] Apart from this extra choice point, the extended structure will be shown to be conservative with respect to the relations of branching and undividedness present in the initial structure.

Our guiding idea is that there should be no correlations in the probability space and among the random variables in the extended structure, but these objects have to *correspond* to the PFB-infested probability space and random variables in the surface structure. We discuss the correspondences between various objects in the surface structure and in the extended structure, defining corresponding transitions, corresponding probability spaces, and corresponding random variables.

The notion of correspondence is needed to formulate a restriction on the function μ' that generates the causal probability spaces in the extended structure. We explain what it means that the propensity assignment μ' in the extended structure adequately represents the propensity assignment μ in the surface structure. Finally, we introduce the BST_{92} notion of a structure with a probabilistic hidden variable for PFB that incorporates two desiderata: (1) the adequate representation of probabilities from the surface structure and (2) no correlations among the hidden variables corresponding to the random variables exhibiting PFB in the surface structure.

To give a roadmap of our formal construction, Def. 8.16 introduces probabilistic BST_{92} surface structures, which form a subclass of probabilistic BST_{92} structures specified in Def. 7.3. Definition 8.18 then explains the notion of an N-multiplied structure corresponding to a probabilistic BST_{92} surface structure. At that stage it is not yet decided whether the latter structure fulfils the axioms of BST_{92}. This issue is decided by Theorem 8.5, which says that an N-multiplied structure corresponding to a probabilistic

[19] Although there is a certain ambiguity in the literature whether to use the plural, "hidden variables", or the singular, we say "values of a (single) hidden variable" since the outcomes of a choice point are representable by a single random variable.

BST_{92} surface structure *is* a BST_{92} structure. The question whether it is a probabilistic BST_{92} structure is still left open. To handle the latter question, we need the notion of an adequate propensity assignment, which is the subject of Def. 8.20. This definition concerns a probabilistic BST_{92} surface structure (with its propensity function μ) and a corresponding N-multiplied BST_{92} structure. The definition spells out what it takes for μ', defined on the transitions of the latter structure, to be *adequate* for the surface structure and N, the size of multiplication. With the notion of adequacy to hand, our next Lemma 8.5 says that an N-multiplied BST_{92} structure corresponding to a surface structure, taken together with a propensity function adequate for the surface structure and for the size N, is a probabilistic BST_{92} structure. At this stage we finally have established the notion of an N-multiplied probabilistic BST_{92} structure corresponding to a probabilistic BST_{92} surface structure. Our construction ends with Def. 8.22, which singles out, from among the N-multiplied probabilistic BST_{92} structures corresponding to a given surface structure with a case of PFB, the *structures with a probabilistic hidden variable for the given case of PFB*.

We will now develop the above-mentioned formal machinery. First, we define the notion of a probabilistic BST_{92} surface structure:[20]

Definition 8.16. A *probabilistic BST_{92} surface structure* is a sextuple $\langle W, <, \mu, e^*, E, C\rangle$, where $\langle W, <, \mu\rangle$ is a probabilistic BST_{92} structure as in Def. 7.3, $e^* \in W$ is a deterministic point in W, and $E, C \subseteq W$ are disjoint sets of choice points that jointly comprise all choice points in $\mathscr{W}$, $e^* < E$ (i.e., for any $e \in E$, we have $e^* < e$), and in W there are only finitely many choice points, and each one is only finitely splitting.

Exactly as in the modal case discussed in Section 8.3, E represents Nature-given indeterminism, whereas C represents indeterminism related to the experimenters' choices of parameters.

We next introduce the auxiliary notion of a lifted history, which we then use to define an N-multiplied structure.

Definition 8.17 (Lifted history). Let $\mathscr{W} = \langle W, <, \mu, e^*, E, C\rangle$ be a probabilistic BST_{92} surface structure and let $h \in \text{Hist}(W)$. The lifted history $\varphi^n(h)$

[20] In contrast to the non-probabilistic surface structures of Def. 8.2, we do not require here that every point in C be below some point in E, as this condition does not play any role in the arguments that follow. The assumption that e^* is deterministic is not essential, but it simplifies the calculations considerably.

for $n \in \mathbb{N}$ is defined as:

$$\varphi^n(h) =_{\mathrm{df}} \{\langle x, 0\rangle \mid x \in h \wedge x \not> e^*\} \cup \{\langle x, n\rangle \mid x \in h \wedge x > e^*\}.$$

Note that if $e^* \notin h$, then $\varphi^n(h) = \{\langle x, 0\rangle \mid x \in h\} = \varphi^m(h)$ for any $m, n \in \mathbb{N}$.

Definition 8.18 (N-multiplied structure)**.** Let $\mathscr{W} = \langle W, <, \mu, e^*, E, C\rangle$ be a probabilistic BST_{92} surface structure. The *N-multiplied structure corresponding to* $\mathscr{W}$ is $\mathscr{W}' = \langle W', <'\rangle$, where

$$W' =_{\mathrm{df}} \bigcup_{h \in \mathrm{Hist}(W), n \in \{1,\dots,N\}} \varphi^n(h),$$

and the ordering $<'$ is given by

$$\langle x_1, n\rangle <' \langle x_2, m\rangle \Leftrightarrow_{\mathrm{df}} x_1 < x_2 \text{ and } (n = m \text{ or } n = 0),$$

where $n, m \in \{0, 1, \dots, N\}$.

We need to show next that an N-multiplied structure is a BST_{92} structure. The first step toward this objective is the following lemma:

Lemma 8.4. *Let $\mathscr{W} = \langle W, <, \mu, e^*, E, C\rangle$ be a probabilistic BST_{92} surface structure and $\mathscr{W}'$—the N-multiplied structure corresponding to $\mathscr{W}$. Then*

(1) *For every history $h \in \mathrm{Hist}(W)$ the set $\varphi^n(h)$ is a maximal directed subset of W', i.e., a history in W'.*
(2) *For every maximal directed subset $A' \subseteq W'$ there is a history $h \in \mathrm{Hist}(W)$ and $n \in \{1, \dots, N\}$ for which $A' = \varphi^n(h)$.*

Proof. See Appendix A.4. □

Our next fact tells us about choice points in the N-multiplied structure corresponding to a probabilistic BST_{92} surface structure.

Fact 8.15. *Let $\mathscr{W} = \langle W, <, \mu, e^*, E, C\rangle$ be a probabilistic BST_{92} surface structure, and let $\mathscr{W}'$ be the N-multiplied structure corresponding to $\mathscr{W}$. Then*

(1) *for every $n \in \{1, \dots, N\}$ and every $h_1, h_2 \in H_{e^*}$: $\varphi^n(h_1) \equiv_{\langle e^*, 0\rangle} \varphi^n(h_2)$.*
(2) *for every $n, m \in \{1, \dots, N\}$ such that $n \neq m$ and every $h \in H_{e^*}$: $\varphi^n(h) \perp_{\langle e^*, 0\rangle} \varphi^m(h)$.*

(3) $\langle e^*, 0\rangle$ *is a choice point with N outcomes* $\Pi_{\langle e^*,0\rangle}\langle\varphi^n(h)\rangle$*, where h is an arbitrary history from* H_{e^*}*;*

(4) *for every* $n \in \{1,\ldots,N\}$*, every* $e \in W$*, and every* $h_1, h_2 \in \mathrm{Hist}(W)$*:* $h_1 \perp_e h_2$ *iff* $\varphi^n(h_1) \perp_{\langle e,l\rangle} \varphi^n(h_2)$*, where* $l = n$ *iff* $e^* < e$*, and* $l = 0$ *otherwise;*

(5) *for every* $m, n, l \in \{1,\ldots,N\}$ *with* $m \neq n$*, every* $e \in W$ *such that* $e^* < e$*, and every* $h_1, h_2 \in \mathrm{Hist}(W)$*: neither* $\varphi^m(h_1) \equiv_{\langle e,l\rangle} \varphi^n(h_2)$*, nor* $\varphi^m(h_1) \perp_{\langle e,l\rangle} \varphi^n(h_2)$*;*

(6) *for every* $m, n \in \{1,\ldots,N\}$ *with* $m \neq n$*, every* $e \not> e^*$*, and every* $h \in H_e$*:* $\varphi^m(h) \equiv_{\langle e,0\rangle} \varphi^n(h)$*.*

Proof. See Appendix A.4. □

The N-multiplied structure adds to the surface structure a choice point $\langle e^*, 0\rangle$ with N outcomes $\Pi_{\langle e^*,0\rangle}\langle\varphi^n(h)\rangle$, where h is any history from H_{e^*}. Once the propensities are added, these outcomes will play the role of values of a hidden variable. Note the difference between clauses (4) and (5)—the former concerns identical superscripts, and the latter is stated for different superscripts. This reflects the fact that any e above e^* is copied into N "new" events $\langle e, n\rangle$, each occurring in a separate history $\varphi^n(h)$, so that it cannot serve as a choice point (nor a point of undividedness) for histories with different superscripts n, m.

To sum up Fact 8.15, as far as choice points are concerned, the N-multiplied structure adds just one single choice point with N outcomes; whatever is above the designated event e^* in the surface structure is multiplied N times: an l-fold choice point is multiplied into N different l-fold choice points, and a deterministic point is multiplied into N deterministic points. In contrast, events that are not above e^* are just copied, not multiplied.

With these results to hand, in full analogy to Theorem 8.1, we can prove that an N-multiplied structure corresponding to a probabilistic BST_{92} surface structure is again a BST_{92} structure.

Theorem 8.5. *Let* $\mathscr{W} = \langle W, <, \mu, e^*, E, C\rangle$ *be a probabilistic* BST_{92} *surface structure, and let* $\mathscr{W}'$ *be the N-multiplied structure corresponding to* $\mathscr{W}$*. Then* $\mathscr{W}'$ *is a* BST_{92} *structure.*

Proof. See Exercise 8.6. □

Note that we have not yet established that an N-multiplied structure corresponding to $\mathscr{W}$ is a *probabilistic* BST_{92} structure. This is essential for analyzing *probabilistic* hidden variables. Before we prove the relevant result, we need to introduce a number of auxiliary notions.

The first notion is that of correspondence. We will advance claims that an N-multiplied structure has no PFB in the region *corresponding* to a PFB-infested region of a surface structure. At the end of the day we will claim that in the N-multiplied structure, transitions and random variables that *correspond*, respectively, to troublesome transitions and troublesome random variables in the surface structure, do not exhibit PFB. Toward this end we need to explain correspondences between various kinds of objects in a probabilistic BST_{92} surface structure and in the extended, N-multiplied probabilistic structure. We begin with transitions:

Definition 8.19 (Corresponding transitions and sets thereof)**.** Let $\mathscr{W} = \langle W, <, \mu, e^*, E, C\rangle$ be a probabilistic BST_{92} surface structure, and let $\mathscr{W}'$ be the N-multiplied structure corresponding to $\mathscr{W}$.

Let $\tau = e \rightarrowtail \Pi_e\langle h\rangle \in \text{TR}(W)$, $T \subseteq \text{TR}(W)$, and $S \subseteq \mathscr{P}(\text{TR}(W))$ be, respectively, a basic transition, a set of basic transitions, and a set of sets of basic transitions in $\mathscr{W}$. The corresponding objects, respectively, τ^n, T^n and S^n in $\mathscr{W}'$ are defined as follows, for $n \in \{1, \ldots, N\}$:

(1) $\tau^n = \langle e, l\rangle \rightarrowtail \Pi_{\langle e,l\rangle}\langle \varphi^n(h)\rangle$, where $l = n$ iff $e^* < e$, and $l = 0$ otherwise;
(2) $T^n = \{\tau^n \mid \tau \in T\}$ and
(3) $S^n = \{T^n \mid T \in S\}$.

As to clause (1), since $\tau = e \rightarrowtail \Pi_e\langle h\rangle$, we have $e \in h$, and hence, with the mentioned caveat about the location of e, $\langle e, l\rangle \in \varphi^n(h)$. Given Fact 8.15(4) it is straightforward to see that τ^n is indeed a basic transition in $\mathscr{W}'$.

Observe that the corresponding basic transitions (and hence sets thereof) are different, depending on the location of their initials. The difference is explained by the following Fact:

Fact 8.16. *Let $\mathscr{W} = \langle W, <, \mu, e^*, E, C\rangle$ be a probabilistic BST_{92} surface structure, and let $\mathscr{W}'$ be the N-multiplied structure corresponding to $\mathscr{W}$. For $e \in W$, $\tau = e \rightarrowtail \Pi_e\langle h\rangle \in \text{TR}(W)$, and $m, n \in \{1, \ldots, N\}$ with $m \neq n$:*

(1) If $e^* < e$, then $\tau^n = \langle e, n\rangle \rightarrowtail \Pi_{\langle e,n\rangle}\langle \varphi^n(h)\rangle \neq \langle e, m\rangle \rightarrowtail \Pi_{\langle e,m\rangle}\langle \varphi^m(h)\rangle = \tau^m$.
(2) If $e^* \not< e$, then $\tau^n = \langle e, 0\rangle \rightarrowtail \Pi_{\langle e,0\rangle}\langle \varphi^n(h)\rangle = \langle e, 0\rangle \rightarrowtail \Pi_{\langle e,0\rangle}\langle \varphi^m(h)\rangle = \tau^m$.

Proof. (1) If $e^* < e$, then e can be only associated with $1, 2, \ldots, N$, and $\langle e, m\rangle \neq \langle e, n\rangle$ for $m \neq n$, so $\tau^m \neq \tau^n$. (2) If $e^* \not< e$, e can be only associated with 0; the claim then follows by Fact 8.15(6). □

Thus, a basic transition $(e \rightarrowtail H) \in \mathrm{TR}(W)$ either has N copies or just one copy in $\mathscr{W}'$, depending on whether $e^* < e$ or not. Note that the clauses exclude the case $e = e^*$, for which a deterministic transition $e^* \rightarrowtail H_{e^*}$ gives rise to N basic transitions $\langle e^*, 0\rangle \rightarrowtail H^n$, where $H^n \in \Pi_{\langle e^*,0\rangle}$.

We work toward assigning propensities to transitions in an N-multiplied structure with a view to turning that structure into a *probabilistic* BST_{92} structure. With the concept of correspondence to hand, we define adequate propensity assignments as follows:

Definition 8.20 (Adequate propensity assignment)**.** Let $\mathscr{W} = \langle W, <, \mu, e^*, E, C\rangle$ be a probabilistic BST_{92} surface structure, and let $\mathscr{W}' = \langle W', <'\rangle$ be the N-multiplied BST_{92} structure corresponding to $\mathscr{W}$. Let

$$E_1 = \{e \in W \mid \mu(\{e \rightarrowtail H\}) \text{ is defined for some } H \in \Pi_e\};$$

by Postulate 7.2, μ is defined on any $Y \subseteq \tilde{T}_{E_1}$, and hence on any element of S_{E_1}. We say that $\mu' : \mathscr{P}(\mathrm{TR}(W')) \mapsto [0,1]$ is *adequate for* $\mathscr{W}$ *and* N iff

(1) μ' is defined for every basic transition $\langle e^*, 0\rangle \rightarrowtail H^n$, where $H^n \in \Pi_{\langle e^*,0\rangle}$;
(2) μ' is defined on each T^n corresponding to some $T \in S_{E_1}$, and
- if there is an initial $e \in E_T$ such that $e^* < e$, then $\mu(T) = \sum_{n=1}^{N} \mu'(\{\langle e^*, 0\rangle \rightarrowtail H^n\}) \cdot \mu'(T^n)$;
- if there is no initial $e \in E_T$ such that $e^* < e$, then $\mu(T) = \mu'(T^n)$;

(3) μ' is defined on every consistent set $\{\langle e^*, 0\rangle \rightarrowtail H^m\} \cup T^n$ (where $m, n \in \{1, \ldots, N\}$), as follows:

$$\mu'(\{\langle e^*, 0\rangle \rightarrowtail H^n\} \cup T^m) = \mu'(\{\langle e^*, 0\rangle \rightarrowtail H^n\}) \cdot \mu'(T^m).$$

The conditions on μ' can be glossed as follows. The first clause requires one to assign values only to the really new transitions (i.e., to transitions having no corresponding transitions in the surface structure). The second clause concerns counterparts of elements of S_{E_1}. Its two parts reflect the difference between corresponding transitions established in Fact 8.16. The first part concerns multiplied transitions, whereas the second concerns non-multiplied transitions (note that in the second part, $T^n = T^m$ for every $m, n < N$). The last clause concerns the mixed case, that is, a set consisting of a new transition and a counterpart of an "old" set (an element of S_{E_1}). In the mixed case, the rule prescribes to multiply the probabilities, as, after all, the indeterministic event $\langle e^*, 0\rangle$ is posited to explain cases of PFB. Allowing for probabilistic dependence between $\langle e^*, 0\rangle$ and a counterpart of the "old" set of transitions would contravene this project. The clause is nevertheless controversial. As we will see, it might encode some independence conditions between C and E, and as we will argue in our analysis of the Bell-Aspect experiment in Section 8.4.4, the condition might actually fail, showing that the deep-structure explanation of Bell-Aspect that we are after is not achievable.

Note that, although the conditions of Def. 8.20 concern elements of S_{E_1}, they induce a propensity assignment to other sets corresponding to subsets of $\tilde{T}_{E_1}$ and their combinations with new transitions as well. That is, for every $Y^n \in \mathrm{TR}(W')$ that corresponds to some $Y \subseteq \tilde{T}_{E_1}$:

- $\mu'(Y^n) = \sum_{T \in S_Y} \mu'(T^n)$, where $S_Y =_{\mathrm{df}} \{T \in S_{E_1} \mid Y \subseteq T\}$, and
- if $\{\langle e^*, 0\rangle \rightarrowtail H^m\} \cup Y^n$ is consistent, then

$$\mu'(\{\langle e^*, 0\rangle \rightarrowtail H^m\} \cup Y^n) = \mu'(\{\langle e^*, 0\rangle \rightarrowtail H^m\}) \cdot \mu'(Y^n).$$

Moreover, a μ' that is adequate for $\mathscr{W}$ and N delivers the surface propensities for any $Y \subseteq \tilde{T}_{E_1}$, not just for elements of S_{E_1}. This can be established as follows:

$$\sum_{n \leqslant N} \mu'(Y^n)\mu'(\{\langle e^*, 0\rangle \rightarrowtail H^n\}) = \sum_{n \leqslant N} \sum_{T \in S_Y} \mu'(T^n) \cdot \mu'(\{\langle e^*, 0\rangle \rightarrowtail H^n\}) =$$
$$\sum_{T \in S_Y} \sum_{n \leqslant N} \mu'(T^n) \cdot \mu'(\{\langle e^*, 0\rangle \rightarrowtail H^n\}) = \sum_{T \in S_Y} \mu(T) = \mu(Y). \tag{8.4}$$

Observe that adequate μ' is sensitive to details of the surface structure: the values of μ, on which sets of transitions μ is defined, and the number N that gives the size of multiplication. All other features of an N-multiplied structure do not bring anything new, as they are copied from the surface structure. This justifies our terminology of μ' being adequate to $\mathscr{W}$ and N.

As expected, the notion of μ' being adequate for $\mathscr{W} = \langle W, <, \mu, E, C\rangle$ and N provides for the following: if one supplements $\mathscr{W}'$ with μ', creating $\mathscr{W}' = \langle W', <', \mu'\rangle$, then the result will be a *probabilistic* BST_{92} structure. This is what our next lemma establishes: a BST_{92} structure corresponding to a probabilistic BST_{92} surface structure is a *probabilistic* BST_{92} structure in the sense of Def. 7.3.

Lemma 8.5. *Let $\mathscr{W} = \langle W, <, \mu, e^*, E, C\rangle$ be a probabilistic BST_{92} surface structure, and let $\mathscr{W}' = \langle W', <'\rangle$ be the N-multiplied BST_{92} structure corresponding to $\mathscr{W}$. Let the partial function $\mu' : \mathscr{P}(\mathrm{TR}(W')) \mapsto [0,1]$ be adequate for $\mathscr{W}$ and N. Then $\mathscr{W}'' =_{\mathrm{df}} \langle W', <', \mu'\rangle$ is a probabilistic BST_{92} structure.*

We call $\mathscr{W}''$ an N-multiplied probabilistic BST_{92} structure corresponding to $\mathscr{W}$.

Proof. See Exercise 8.7. □

The requirement of an adequate propensity assignment is part of the notion of an N-multiplied probabilistic BST_{92} structure corresponding to $\mathscr{W}$. Since the requirement is fulfilled by many (partial) functions μ' that are adequate for a given surface probabilistic structure $\mathscr{W}$ and N, there are many N-multiplied probabilistic BST_{92} structures corresponding to $\mathscr{W}$. However, they all recapture the propensities of the surface structure, due to clause (2) of Def. 8.20.

At this stage we have defined probabilistic BST_{92} surface structures and the N-multiplied probabilistic BST_{92} structures corresponding to them. We now define the remaining correspondences we need, first between probability spaces and then between random variables.

Definition 8.21 (Corresponding probability spaces and corresponding random variables). Let $\mathscr{W} = \langle W, <, \mu, e^*, E, C\rangle$ be a probabilistic BST_{92} surface structure, and let $\mathscr{W}' = \langle W', <', \mu'\rangle$ be an N-multiplied probabilistic BST_{92} structure corresponding to $\mathscr{W}$. Let $CPS = \langle S, \mathscr{A}, p\rangle$ be a causal probability space, with $S \subseteq \mathrm{TR}(W)$.

Then a triple $CPS^n = \langle S^n, \mathscr{A}^n, p^n\rangle$ with $n \leqslant N$ is *the n-th causal probability space corresponding to CPS* iff S^n and $\mathscr{A}^n$ correspond to S and $\mathscr{A}$, respectively (in the sense of Def. 8.19), and p^n is induced by the propensity assignment μ' ($n \leqslant N$).

Further, for random variables X and X^n defined within the corresponding probability spaces CPS and CPS^n, respectively, we say that *X and X^n correspond* iff X and X^n have the same range and

$$\text{for every } T \in S: \quad X(T) = x \text{ whenever } X^n(T^n) = x,$$

where $T^n \in S^n$ is the set corresponding to $T \in S$.

Note that, typically, corresponding random variables have different probabilities, i.e., $p(X(T) = x) \neq p^n(X^n(T^n) = x)$, because typically μ assigns to T a different propensity than μ' assigns to T^n.

Having all the machinery in place, we can now apply it to a surface structure with a case of PFB. To this end let us first recall our description of PFB, as offered by Def. 7.7:

Given a probabilistic BST_{92} structure, a locus of PFB is provided by pairwise SLR initial events $I_1, \ldots, I_K$, each associated with a transition to an unavoidable disjunctive outcome $\mathbf{1}_k$ with cardinality $\Gamma(k)$. The set of transitions $\{I_1 \rightarrowtail \mathbf{1}_1, \ldots, I_K \rightarrowtail \mathbf{1}_K\}$ is said to exhibit PFB iff the random variables $\{X_1, \ldots, X_K\}$ are correlated. These random variables are defined on the causal probability space $CPS(E \rightarrowtail \mathbf{1}_E) = \langle S, \mathscr{A}, p \rangle$ determined by the transition $E \rightarrowtail \mathbf{1}_E$, where $E = \bigcup_{k=1}^{K} I_k$ and $\mathbf{1}_E = \{\bigcup Z \mid Z \in \mathbf{1}_1 \times \mathbf{1}_2 \times \ldots \times \mathbf{1}_K\}$, by the formula:

$$X_k : S \mapsto \Gamma(k) \text{ where for every } T \in S : X_k(T) = \gamma \text{ iff } CC(I_k \rightarrowtail \hat{O}_\gamma) \subseteq T. \tag{8.5}$$

Having recalled the BST_{92} analysis of PFB, we end our construction by singling out, from among the N-multiplied probabilistic BST_{92} structures corresponding to a given surface structure that includes a case of PFB, a *structure with a probabilistic hidden variable for the given case of PFB*.

Definition 8.22 (Structure with a probabilistic hidden variable for PFB)**.** Let $\mathscr{W} = \langle W, <, \mu, e^*, E, C \rangle$ be a probabilistic BST_{92} surface structure with transitions $\{I_1 \rightarrowtail \mathbf{1}_1, \ldots, I_K \rightarrowtail \mathbf{1}_K\}$, where $\mathbf{1}_k = \{\hat{O}_{k,\gamma(k)} \mid \gamma(k) \in \Gamma(k)\}$, and random variables $X_1, \ldots, X_K$ of Def. 8.5 that exhibit PFB. Let also $cll(I_k \rightarrowtail \hat{O}_{k,\gamma(k)}) \subseteq E$, for every $k \leqslant K$ and every $\gamma(k) \in \Gamma(k)$.
A *structure with a probabilistic hidden variable for the given PFB in* $\mathscr{W}$ is an N-multiplied probabilistic BST_{92} structure $\mathscr{W}' = \langle W', <', \mu' \rangle$ (for some $N \in \mathbb{N}$) that corresponds to $\mathscr{W}$ and which satisfies the following condition:
for every $n \leqslant N$, the random variables $X_1^n, \ldots, X_K^n$ corresponding to $X_1, \ldots, X_K$ are independent. [Outcome Independence]

To see that our condition captures the condition of Outcome Independence as it is known from the literature, observe that each random variable $X_k : S \mapsto \Gamma(k)$ represents the results of a measurement I_k; a corresponding random variable X_k^n represents the counterparts of these results in the N-multiplied structure for the n-th value of the hidden variable. The initials $I_1, \ldots, I_K$ are pairwise *SLR*, and so are their images. The statement that the latter random variables are independent means that the probabilities of the images of results in the N-multiplied structure (for each $n \leqslant N$) are independent (i.e., they multiply to yield the probability of the joint result). Observe that the defined notion incorporates two conditions, Outcome Independence and Adequate Propensity Assignment.

Note that the general procedure of N-multiplication can introduce a case of PFB that has no counterpart in a surface structure. This is connected with the fact that the factorization of propensities in the surface structure does not guarantee the factorization of propensities in an N-multiplied structure. That is, $\mu(T_1 \cup T_2) = \mu(T_1) \cdot \mu(T_2)$ does not imply $\mu'(T_1^n \cup T_2^n) = \mu'(T_1^n) \cdot \mu'(T_2^n)$, where T_i and T_i^n are a set of transitions in the surface structure and a set corresponding to it in an N-multiplied structure, respectively. The adequacy condition only requires that the μ-propensity results as the weighed average of the μ' propensities. Whether the notion of N-multiplied BST_{92} probabilistic structure corresponding to a surface structure should prohibit such new cases of PFB, depends on one's attitude as to what the notion is to achieve. If its aim is to *remove* the PFB, the prohibition is very much in place. But if its aim is to *explain* the PFB in the surface structure, the prohibition should not necessarily be imposed, as there might be no way to account for the surface PFB apart from also acknowledging PFB on a deeper level. In what follows, we do not impose the mentioned restriction.

Having explained our framework for the construction of a probabilistic hidden variable for PFB, we finally turn to interesting real questions: Do various set-ups with cases of PFB, of varying complexity, admit structures with probabilistic hidden variables for their cases of PFB?

8.4.3 Single and multiple cases of PFB, and super-independence

Before we investigate BST_{92} structures with PFB, which come with specified sets of experimenter-controlled choice points C and Nature's choice points E, we need to add more substance to the notion of C/E independence.

We introduced C/E independence as a target notion in Def. 8.1, and we already gave it a rigorous formulation in the context of deterministic hidden variables (see Def. 8.3 for surface structures and Def. 8.10 for extended structures). We now have to flesh out the notion in the context of probabilistic hidden variables, by relating it to propensities. For this approach to go through, however, we need to assign propensities to certain sets of transitions that include both an agent-based transition and a Nature-given transition. This move may be controversial, so we try to keep our commitments to propensities of agent-based transitions minimal. In our view, an unproblematic case is a transition $I \rightarrowtail \hat{O}$ from before an experimenter's making a measurement decision to the occurrence a measurement result. The set $CC(I \rightarrowtail \hat{O})$ contains both the experimenter-based transition and a Nature-given transitions. In such a case, when a set T with "mixed" transitions is identical to the set of *causae causantes* of some transition, it may (we think) have a propensity assigned. Once the propensity is assigned to such a set, the notion C/E propensity independence is easily formalized, by equating it with the factorization of propensities. (Note that by Postulate 7.2, if a propensity is defined on a set, it is defined on any of its subsets.) These ideas are reflected by the following definition:

Definition 8.23 (C/E propensity independence)**.** Let $\mathscr{W} = \langle W, <, \mu, E, C\rangle$ be a probabilistic BST$_{92}$ structure with two designated disjoint sets E, C of choice sets such that $E \cup C$ is the set of all choice points in $\mathscr{W}$. We say that $\mathscr{W}$ *violates C/E propensity independence* iff there is a transition Tr in W and $T = CC(Tr)$ with $cll(T) \cap C = C_0 \neq \emptyset$ and $cll(T) \cap E = E_0 \neq \emptyset$ such that $\mu(T) \neq \mu(T_{C_0}) \cdot \mu(T_{E_0})$, where T_{C_0} is the subset of T with initials in C_0, and T_{E_0} is the subset of T with initials in E_0, i.e., $T_{C_0} =_{\text{df}} \{(e \rightarrowtail H) \in T \mid e \in C_0\}$ and $T_{E_0} =_{\text{df}} \{(e \rightarrowtail H) \in T \mid e \in E_0\}$.

We say that a probabilistic BST$_{92}$ structure (or its set E) *satisfies* C/E independence iff it does not violate it.

In what follows we apply this definition to both probabilistic surface structures and to N-multiplied probabilistic structures; for this reason we have not included e^* in the specification of $\mathscr{W}$.

8.4.3.1 A structure with a single case of PFB

We first investigate a simple set-up that exhibits a single case of PFB and which does not involve any relevant choices of experimenters. We produce a probabilistic BST$_{92}$ surface structure for this set-up and ask if it can be given an N-multiplied BST$_{92}$ probabilistic structure with a probabilistic hidden variable for this case of PFB. The absence of relevant choices of experimenters means that $C = \emptyset$ in the description of the surface structure for this set-up. The following lemma proves that the answer to our question is "yes", given that some minor conditions on the surface structure are satisfied.

Lemma 8.6. *Let* $\mathscr{W} = \langle W, <, \mu, e^*, E, C\rangle$ *be a probabilistic BST$_{92}$ surface structure in which transitions* $\{I_1 \rightarrowtail \mathbf{1}_1, \ldots, I_K \rightarrowtail \mathbf{1}_K\}$ *with random variables* $X_1, \ldots, X_K$ *exhibit PFB, where* $\mathbf{1}_k = \{\hat{O}_{k,\gamma(k)} \mid \gamma(k) \in \Gamma(k)\}$, $\Gamma(k)$ *are index sets and* $1 \leqslant k \leqslant K$. *Let* $cll(I_k \rightarrowtail \hat{O}_{k,\gamma(k)}) \subseteq E$ *for every* $k \leqslant K$ *and every* $\gamma(k) \in \Gamma(k)$ *and* $C = \emptyset$. *Then there exists a structure with a probabilistic hidden variable for this case of PFB. Moreover, the structure satisfies* C/E *propensity independence.*

Proof. See Appendix A.4. □

This lemma says that if we focus on a single case of PFB, if only there is a deterministic event below all cause-like loci involved in this case of PFB and our common finitistic assumptions are satisfied, then there is a probabilistic BST$_{92}$ structure that explains away the PFB in question. The result holds no matter what the background of the case of PFB consists of. In the background there might be some other transitions, possibly forming other cases of PFB. And, on some sets of such neglected transitions, μ might be undefined in the surface structure. Still, we can ignore all such complexities, the lemma says, and construct a probabilistic BST$_{92}$ structure explaining away a single case of PFB. The fact that μ might fail to be defined on a set of transitions has an impact, however, on the feasibility of a more general result concerning multiple cases of PFB.[21] Accordingly, in our next result, we assume that μ is defined on all consistent subsets of the set of all basic transitions.

8.4.3.2 A structure with multiple cases of PFB — super-independence

Given the success of the construction of a BST$_{92}$ structure with a probabilistic hidden variable for a set-up with a single instance of PFB and no choices produced by agents, it is natural to ask whether the construction can be successfully applied to set-ups with two or more instances of PFB. This category includes set-ups with and without agent-based selections of parameters. We can prove that all set-ups with finitely many cases of PFB and

[21] For starters, attempt to analyze two cases of PFB, one "above" the other, with some transitions "between" the two, that do not form a set to which μ is assigned.

no agent-based choices admit an N-multiplied BST_{92} probabilistic structure in which the instances of PFB are removed. The sought-after N-multiplied structure can be produced in a way that is quite similar to construction of the quasi-deterministic structure investigated in the proof of Lemma 8.6. The difference is that the present lemma concerns *all* basic transitions in the surface structure, in contrast to just those on which μ is defined. The relevant lemma is as follows.

Lemma 8.7. *Let $\mathscr{W} = \langle W, <, \mu, e^*, E, C\rangle$ be a probabilistic BST_{92} surface structure harboring multiple cases of PFB for which $C = \emptyset$. Let μ be defined on every subset of* $\mathrm{TR}(W)$*. Then there is an N-multiplied probabilistic BST_{92} structure corresponding to $\mathscr{W}$ that provides a hidden variable for every case of PFB in $\mathscr{W}$. Moreover, that extended structure satisfies C/E propensity independence.*

Proof. See Appendix A.4. □

We will call the option of removing *all* cases of PFB, without paying attention to other constraints, "super-independence", in analogy to "super-determinism". Superdeterminism, as investigated in Section 8.3.1.1, is an option for removing all the (finitely many) cases of MFB in a BST_{92} surface structure in which $C = \emptyset$. Theorem 8.2 shows that if our single aim is to remove cases of surface MFB by positing instruction sets, we can always achieve this by assuming that the experimenters' choices are in fact due to Nature ($C = \emptyset$). Similarly, if our single aim is to explain away multiple cases of PFB, we can always achieve this—we just need to assume that no indeterminism comes from experimenters' choices, so that $C = \emptyset$. In this case, any constraints related to members of C are satisfied vacuously. Lemma 8.7 guarantees that such super-independent extensions always exist.

As we said repeatedly, the challenge in constructing hidden variable extensions of surface structures consists in analyzing the set-ups in question *as experiments*; that is, in a way that upholds a separation between Nature's and experimenters' choices ($C \neq \emptyset$). We now turn to the Bell-Aspect experiment,[22] in which experimenters' choices *are* present. It is generally assumed that a good explanation of this experiment should accommodate such choices, and superdeterministic or super-independent accounts are dismissed as conspiratorial. We will show that the existence of an extended

[22] The set-up was proposed in Bell (1964), and a breakthrough experiment with the set-up was carried out by Aspect et al. (1982b).

structure with a hidden variable, given the cases of PFB present in this experiment, implies the so-called Bell-CH inequality. As the Bell-CH inequality is violated by quantum mechanical predictions and in many experiments, we arrive at a "no go" result concerning a propensity-based account of non-local correlations. Our formal derivation of the Bell-CH inequality is standard,[23] but our construction, which focuses on both the modal (propensity) and spatio-temporal aspects, is novel. The thrust of our analysis is an attempt to understand, having support of all the resources of BST, what the premisses for the derivation of the Bell-CH inequality amount to. In this way we hope to contribute to the answering of an ultimately metaphysical question: What must our world be like for the Bell-CH inequality to fail?

8.4.4 The Bell-Aspect experiment

In the last section we developed the technique of structure multiplication. A multiplied surface structure is intended to explain away instances of PFB present in the surface structure. If successful, the construction permits one to interpret PFB epistemically, by claiming that on a deeper level, as described by the N-multiplied structure, there is no PFB. PFB can be seen only on a less than fully fine-grained description of the phenomena; it comes from averaging over the deep level probabilities. We have seen some successes with this program: We proved that any single case of PFB can be explained away by N-multiplication. Also, we showed that, if we ignore agent-based choices as a separate category (subsuming them under the category of Nature-given indeterminism as well), any finite number of cases of PFB can be explained away by the same method. Given the unproblematic development of the technique and its initial successes, the rules of suspense suggest that one should now expect a plot twist. The literature on mysterious quantum correlations additionally enhances the feeling of an imminent catastrophe: an inexplicable set-up with PFB is around the corner.

These feelings are fully justified. In this section we analyze an experiment, the famous Bell-Aspect experiment, for which our N-multiplication technique fails. We will prove, following similar arguments in the literature, that there is no BST_{92} structure with a probabilistic hidden variable for the

[23] Our derivation is a BST rendition of J. S. Bell's reasoning sketched in the introduction to his book (Bell, 1987a).

cases of PFB present in this experiment. Acknowledging that our argument is standard may raise doubts as to why to run it again. The argument has a certain set of premises, so presenting it anew may show that some new premises are needed, or, to the contrary, that some standard premises are superfluous. We will, however, not present any such discoveries about premises of this argument. We rerun the argument to see what it takes, in terms of the modal, causal, propensity-like and spatio-temporal features of our world, for the derivation to go through—or, to put it more technically, to see what it takes for the argument's premises to hold. BST is a suitable framework to address these questions, as it offers all the resources needed in an integrated framework and, importantly, it is mathematically rigorous. Thus, it seem well-suited to address the question that Jeremy Butterfield (1992) raised almost three decades ago: "Bell's theorem: what does it take?"

Our plan is thus to run in the BST framework a standard argument for a "no go" theorem for the existence of local hidden variables for the Bell-Aspect set-up, relying on the Bell-CH inequality. We will identify the premises in terms of BST notions, and, using BST resources, we will investigate what these premises amount to. The aim is to shed some light on what the theorem intimates about our world. What must our locally indeterministic and relativistic world be like for the Bell-CH inequality to fail?

8.4.4.1 The set-up of the Bell-Aspect experiment

The set-up of the Bell-Aspect experiment is outlined in Figure 8.1 (p. 226). Here are some more details. A source emits pairs of particles with spin $\frac{1}{2}$, with each pair being in the singlet spin state already mentioned in Chapter 5, written in the basis $|\pm,\pm\rangle = |\pm\rangle_1 \otimes |\pm\rangle_2$ as

$$|\psi\rangle = \frac{1}{\sqrt{2}}(|+,-\rangle - |-,+\rangle).$$

The members of each pair fly in opposite directions toward remote measurement stations (wo)manned by Alice on the left and Bob on the right. For each emission, in the left station there is Alice's selection a of one of the two settings, a_1 or a_2, of her measuring apparatus, and in the right station there is Bob's selection b of one of the two settings, b_3 or b_4, of his measuring apparatus. The measurement with the setting a_i has two possible results, $a_{i,+}$ or $a_{i,-}$; the measurement with the setting b_j has also two possible results, $b_{j,+}$ or $b_{j,-}$. To write quantified formulas, we henceforth assume that the range of i, i' is $\{1,2\}$, the range of j, j' is $\{3,4\}$ and the range of m, m' is $\{+,-\}$. The relevant *SLR* relations are as follows. In each round of the experiment, each selection event is *SLR* to the measurement event in the remote station (i.e., a *SLR* b_j and b *SLR* a_i). Next, outcomes in different stations are *SLR* (i.e., $a_{i,m}$ *SLR* $b_{j,n}$). Also, a measurement event in one station is *SLR* to an outcome of a measurement event in a remote station (i.e., a_i *SLR* $b_{j,m'}$ and b_j *SLR* $a_{i,m}$). Finally, there are non-local correlations: for each pair (i, j), the remote results $a_{i,m}$ and $b_{j,m'}$ are probabilistically correlated. Such non-local correlations, if understood as reflecting underlying propensities, provide evidence for PFB. The Bell-Aspect experiment involves two kinds of chanciness (like the GHZ experiment): experimenter-induced indeterminism in transitions $a \rightarrowtail a_i$ and $b \rightarrowtail b_j$, and Nature-given chanciness to be seen in transitions $a_i \rightarrowtail a_{i,m}$ and $b_j \rightarrowtail b_{j,m'}$.

8.4.4.2 The surface structure for the Bell-Aspect experiment

We turn now to the construction of a probabilistic BST$_{92}$ surface structure representing the Bell-Aspect experiment, $\mathscr{W}_{BA} = \langle W, <, \mu, e^*, E, C\rangle$. W must contain at least 15 events, which we assume to be all point-like. These are: $a, b, a_i, b_j, a_{i,m},\ b_{j,m'}, e^* \in W$. The ordering relations, including *SLR* relations and compatibility, are specified as follows: $a < a_i$, $b < b_j$, $a_i < a_{i,m}$, $b_j < b_{j,m'}$, then $\Pi_a\langle a_1\rangle \neq \Pi_a\langle a_2\rangle$, $\Pi_b\langle b_3\rangle \neq \Pi_b\langle b_4\rangle$, $\Pi_{a_i}\langle a_{i,+}\rangle \neq \Pi_{a_i}\langle a_{i,-}\rangle$, $\Pi_{b_j}\langle b_{j,+}\rangle \neq \Pi_{b_j}\langle b_{j,-}\rangle$. Next, a *SLR* b_j, b *SLR* a_i, $a_{i,m}$ *SLR* $b_{j,m'}$, and a_i *SLR* $b_{j,m'}$ and b_j *SLR* $a_{i,m}$. And e^* is *SLR* or below each of a, b but $e^* < a_i$ and $e^* < b_j$. It is easy to calculate that $\mathscr{W}_{BA}$ has 16 histories that can be identified via outcomes $a_{i,m} \cup b_{j,m'}$. The a and b-based transitions are agent-induced, whereas the a_i and b_j-based transitions are thought of as Nature-given. Accordingly, $E = \{a_1, a_2, b_3, b_4\}$ and $C = \{a, b\}$. We may further assume that the structure satisfies C/E propensity independence.[24] The next element, the propensity function μ, is assigned only to those objects for which quantum mechanics offers numerical predictions. In this vein, $\mu(a_i \cup b_j \rightarrowtail a_{i,m} \cup b_{j,m'})$ is identified with the QM probability for the joint

[24] For this assumption we need to define propensity on transitions based on subsets of C, however.

outcome $(a_{i,m}, b_{j,m'})$, and $\mu(a_i \rightarrowtail a_{i,m})$ is the QM probability for a single outcome $a_{i,m}$, and analogously for transitions based on b_j.[25]

The ordering relations and the splitting of histories allow us to speak of a number of transitions: $a \rightarrowtail a_i$, $b \rightarrowtail b_j$, $a_i \rightarrowtail a_{i,m}$ and $b_j \rightarrowtail b_{j,m'}$. Although this notation clearly suggests event-like transitions, we use it freely to refer to the corresponding proposition-like transition, canonically written as $a \rightarrowtail \Pi_a\langle a_i\rangle$, $b \rightarrowtail \Pi_b\langle b_j\rangle$, etc. The assumed idealizations imply that all these transitions are basic, so for each transition its set of *causae causantes* consists only of the set in question (i.e., we have identities like $CC(a \rightarrowtail a_i) = \{a \rightarrowtail a_i\}$).

Note that by constructing a surface probabilistic structure in accordance with Def. 8.16, we decide to have just one point event e^* to take care of four cases of PFB present in the set-up. An alternative idea is to postulate four such points, one for each case of PFB. We argue in App. A.4.1 that the option with four events either reduces to our option with one event, or it gives up on explaining PFB.

8.4.4.3 Probabilistic funny business

Our surface structure contains four cases of PFB, and we have to decide whether to follow a "big space" approach, or, alternatively, a "small spaces" approach (see Butterfield 1992). A "big space" approach takes the selection events a and b to construct one "big" causal probability space based on transitions to disjunctive outcomes, $a \cup b \rightarrowtail \{a_i \cup b_j \mid i \in \{1,2\}, j \in \{3,4\}\}$. This approach has a certain mathematical elegance, but disturbingly it assigns propensities to sets like $a \rightarrowtail a_1$. Ideologically, one might oppose to assigning numerical propensities to agent's choices, but even if one does not oppose to such an assignment on ideological grounds, it is unclear where such numbers could come from. We are thus after a "small spaces" approach, which constructs four separate causal probability spaces, each based on transitions $S_{ij} = \{a_i \cup b_j \rightarrowtail \{a_{i,m} \cup b_{j,m'} \mid m, m' \in \{-,+\}\}\}$ (for some fixed i and j). The algebra $\mathscr{A}_{ij}$ of subsets of S_{ij} is induced automatically, and the probabilities

[25] As the settings stand for directions of spin projection, the quantum mechanical probability for a joint measurement on a pair in the singlet state is

$$p_{qm}(a_{i,+}, b_{j,+}) = \frac{1}{2}\cos^2\left(\frac{\angle(i,j)}{2}\right),$$

which means that there are no correlations only if the angle $\angle(i,j)$ between polarization directions equals $\pi/2$. Angles in a typical Bell-type experiment are $\angle(1,3) = \frac{2}{3}\pi = \angle(2,4)$, $\angle(1,4) = \frac{4}{3}\pi$ and $\angle(2,3) = 0$. Thus, for each of these angles there is PFB.

p_{ij} are induced by μ, the values of which are in turn dictated by quantum mechanics. Importantly, for small probability spaces, we need no other values of probabilities than those that are ascribed by quantum mechanics to joint and single measurement results. In this "small spaces" approach, experimenters' choices are out of the picture, as transitions like $a \rightarrowtail a_i$ are not in the algebra of subsets of S_{ij}. (Such transitions will nevertheless return in our discussion, when we attempt to justify the assumptions of Bell's theorem or discuss C/E independence.) The result of this construction is that each case of PFB present in the Bell-Aspect set-up is analyzed in a different probability space.

Given a causal probability space $CPS_{ij} = \langle S_{ij}, \mathscr{A}_{ij}, p_{ij} \rangle$, a case of PFB is specified by transitions $a_i \rightarrowtail \mathbf{1}_{a_i}$ and $b_j \rightarrowtail \mathbf{1}_{b_j}$, where $\mathbf{1}_{a_i} = \{a_{i,+}, a_{i,-}\}$ and $\mathbf{1}_{b_j} = \{b_{j,+}, b_{j,-}\}$. As for the associated random variables, there is a certain subtlety in our notation, since we need to indicate the probability space on which they are defined so as to distinguish X_i as defined on S_{ij} and X_i as defined of $S_{ij'}$. So we use extended subscripts, writing $X_{i,ij}$, to indicate that the random variable in question is defined on the causal probability space CPS_{ij}. With this little notational complication, we assume that every associated random variable has the same range $\Gamma = \{+, -\}$, and we define it in accordance with Def. 7.7: For any $T \in S_{ij}$,

$$X_{i,ij}(T) = m \text{ iff } a_i \rightarrowtail a_{i,m} \in T \quad \text{and} \quad X_{j,ij}(T) = m' \text{ iff } b_j \rightarrowtail b_{j,m'} \in T. \quad (8.6)$$

By the quantum mechanical probabilities for joint outcomes in the Bell-Aspect experiment, the random variables $X_{i,ij}$ and $X_{j,ij}$ are dependent. Thus, $a_i \rightarrowtail \mathbf{1}_{a_i}$ and $b_j \rightarrowtail \mathbf{1}_{b_j}$, together with the associated random variables $X_{i,ij}$ and $X_{j,ij}$, exhibit PFB.

8.4.4.4 Derivation of the Bell-CH inequality

We now turn to the crucial question: is there a BST_{92} structure $\mathscr{W}'_{BA}$ with a probabilistic hidden variable for all four cases of PFB present in the surface structure $\mathscr{W}_{BA}$? We will run a standard derivation of the Bell-CH inequality[26] within the BST framework, intended to show that the answer is in the negative. To recall the logic of the argument, it attempts to show that if the mentioned structure exists, then the Bell-CH inequality has to hold. However, quantum mechanical predictions, supported by overwhelming

[26] See Clauser and Horne (1974); Myrvold et al. (2019).

experimental evidence, show that the inequality is violated. Hence the sought-after structure with a hidden variable for the four cases of PFB does not exist. As our main task is the analysis of the premises of the derivation from the perspective offered by BST, we present the derivation without much ado, to focus later on the justification of its premises.

As to the premises, apart from the conditions of Adequate Propensity Assignment and Outcome Independence that are part and parcel of the definition of a structure with a probabilistic hidden variable for PFB, we assume a constraint known as Parameter Independence, which we define as follows:

Definition 8.24 (Parameter Independence)**.** Let CPS^n_{ij}, $CPS^n_{ij'}$, and $CPS^n_{i'j}$ be causal probability spaces corresponding to spaces CPS_{ij}, $CPS_{ij'}$, and $CPS_{i'j}$, respectively, and $X^n_{i,\,ij}$, $X^n_{i',\,i'j}$, and let $X^n_{j',\,ij'}$ be random variables corresponding to random variables defined by Eq. (8.6). Then we say that $X^n_{i,\,ij}$ and $X^n_{i,\,ij'}$ *satisfy Parameter Independence* iff for every m: $p^n_{ij}(X^n_{i,\,ij} = m) = p^n_{ij'}(X^n_{i,\,ij'} = m)$. Analogously, $X^n_{j,\,ij}$ and $X^n_{j,\,i'j}$ *satisfy Parameter Independence* iff for every m: $p^n_{ij}(X^n_{j,\,ij} = m) \ = p^n_{i'j}(X^n_{j,\,i'j} = m)$.

We now assume that there is a BST_{92} structure $\mathscr{W}'_{BA}$ with a probabilistic hidden variable for all four cases of PFB present in the surface structure $\mathscr{W}_{BA}$ and satisfying Parameter Independence. In more detail, $\mathscr{W}'_{BA}$ is an N-multiplied probabilistic BST_{92} structure that corresponds to the probabilistic BST_{92} surface structure $\mathscr{W}_{BA}$. Our derivation starts from the arithmetical fact that for any real numbers u, u', v, and v' from the unit interval $[0, 1]$,

$$-1 \leqslant uv + uv' + u'v' - u'v - u - v' \leqslant 0. \tag{8.7}$$

As values of probabilities fall into the unit interval, we next make these substitutions:

$$u = p^n_{13}(X^n_{1,13} = +) \quad u' = p^n_{23}(X^n_{2,23} = +) \tag{8.8}$$

$$v = p^n_{13}(X^n_{3,13} = +) \quad v' = p^n_{14}(X^n_{4,14} = +). \tag{8.9}$$

Then, using Outcome Independence and Parameter Independence a few times, we arrive at the following:

$$uv = p^n_{13}((X^n_{1,13} = +) \wedge (X^n_{3,13} = +)) \tag{8.10}$$

$$uv' = p^n_{13}(X^n_{1,13} = +)p^n_{14}(X^n_{4,14} = +) = p^n_{14}(X^n_{1,14} = +)p^n_{14}(X^n_{4,14} = +) = p^n_{14}((X^n_{1,14} = +) \wedge (X^n_{4,14} = +)) \tag{8.11}$$

$$u'v' = p^n_{23}(X^n_{2,23} = +)p^n_{14}(X^n_{4,14} = +) = p^n_{24}(X^n_{2,24} = +)p^n_{24}(X^n_{4,24} = +) = p^n_{24}((X^n_{2,24} = +) \wedge (X^n_{4,24} = +)) \tag{8.12}$$

$$u'v = p^n_{23}((X^n_{2,23} = +) \wedge (X^n_{3,23} = +)). \tag{8.13}$$

Finally, after making the substitutions (8.8)–(8.9) and then processing the transformations (8.10)–(8.13) in inequality (8.7), we multiply the resulting formula side-ways by $\mu'(\langle e^*, 0\rangle \rightarrowtail H^n)$, and then sum over $n \in \{1, \dots, N\}$. Since μ' is a propensity function, $\sum_{n=1}^{N} \mu'(\langle e^*, 0\rangle \rightarrowtail H^n) = 1$. By the condition of Adequate Propensity Assignment (Def. 8.20), we have these identities:

$$\sum_{n \leqslant N} \mu'(\langle e^*, 0\rangle \rightarrowtail H^n) \cdot p^n_{ij}((X^n_{i,ij} = +) \wedge (X^n_{j,ij} = +)) = p_{ij}((X_{i,ij} = +) \wedge (X_{j,ij} = +)). \tag{8.14}$$

Putting these observations together, we arrive at the Bell-CH inequality:

$$-1 \leqslant p_{13}((X_{1,13} = +) \wedge (X_{3,13} = +)) + p_{14}((X_{1,14} = +) \wedge (X_{4,14} = +)) + p_{24}((X_{2,24} = +) \wedge (X_{4,24} = +)) - p_{23}((X_{2,23} = +) \wedge (X_{3,23} = +)) - p_{13}(X_{1,13} = +) - p_{14}(X_{4,14} = +) \leqslant 0. \tag{8.15}$$

A reader familiar with Bell's theorems might ask where we used the assumption that the values of a hidden variable and the measurement settings are independent (i.e., the so-called No Conspiracy assumption). This assumption is at work in Eq. (8.14), which says that there is a settings-independent probability distribution on outcomes of $\langle e^*, 0\rangle$. (Note that settings-dependence of μ' would block the derivation.) No Conspiracy, in one formulation, says that "[t]he probability distribution μ' of [a hidden variable] should not be allowed to depend on (a_i, b_j); this is the mathematical meaning of the assumption… that the control parameters a_i, b_j are "randomly and freely chosen by the experimenters" " (Goldstein et al., 2011).[27] The assumptions of No Conspiracy and Adequate Propensity Assignment point to problem involving a causal discrepancy in the analysis, which we are about to uncover. We proceed with the premises of the derivation, asking the question of what the world must be like in order to justify the premises and the subsequent steps in the above derivation.

8.4.4.5 Analysis of the derivation

As a way toward a BST_{92} analysis of the derivation, recall first that, by the definition of surface probabilistic structures, e^* is below every a_i and b_j. Consider then the causal situation, first in the surface structure $\mathscr{W}_{BA}$, and focus on $CC(e^* \rightarrowtail a_{i,m} \cup b_{j,m'})$. Given the location of e^*, this set must contain $a_i \rightarrowtail a_{i,m}$ and $b_j \rightarrowtail b_{j,m'}$. But does $CC(e^* \rightarrowtail a_{i,m} \cup b_{j,m'})$ contain other transitions as well? Clearly, e^* is not, and cannot be, above a; otherwise, since a is a choice point for measurement outcomes, e^* would prohibit the occurrence of some outcome of a. For an analogous reason, e^* cannot be above b. It follows that e^* is *SLR* or below each a and b. Hence $a \rightarrowtail a_i$ and $b \rightarrowtail b_j$ must belong to $CC(e^* \rightarrowtail a_{i,m} \cup b_{j,m'})$, i.e.,

$$CC(e^* \rightarrowtail a_{i,m} \cup b_{j,m'}) = \{a_i \rightarrowtail a_{i,m}, b_j \rightarrowtail b_{j,m'}, a \rightarrowtail a_i, b \rightarrowtail b_j\}. \quad (8.16)$$

We turn next to the transition $\langle e^*, 0\rangle \rightarrowtail (a_{i,m} \cup b_{j,m'})^n$ in $\mathscr{W}'_{BA}$, where $(a_{i,m} \cup b_{j,m'})^n$ is the n-th counterpart in $\mathscr{W}'_{BA}$ of the event $a_{i,m} \cup b_{j,m'}$. Its set of *causae causantes* almost mirrors the above set of transitions, the difference being the inclusion of $\langle e^*, 0\rangle \rightarrowtail H^n$, where H^n is the n-th elementary outcome of $\langle e^*, 0\rangle$. To write down the set:

$$CC(\langle e^*, 0\rangle \rightarrowtail (a_{i,m} \cup b_{j,m})^n) =$$
$$\{\langle e^*, 0\rangle \rightarrowtail H^n,\ (a_i \rightarrowtail a_{i,m})^n,\ (b_j \rightarrowtail b_{j,m'})^n,\ (a \rightarrowtail a_i)^n,\ (b \rightarrowtail b_j)^n\}. \quad (8.17)$$

Clearly, this is different from the set

$$CC^* = \{\langle e^*, 0\rangle \rightarrowtail H^n,\ (a_i \rightarrowtail a_{i,m})^n,\ (b_j \rightarrowtail b_{j,m'})^n\}. \quad (8.18)$$

[27] We adapted the symbolism in this quote to match the present text.

The difference lies precisely in the C-based (agent-induced) transitions $(a \rightarrowtail a_i)^n$ and $(b \rightarrowtail b_j)^n$, which are members of the first set but not of the second one.

Now, returning to the condition that the N-multiplied structure must reflect the propensities of the surface structure, the underlying idea is that a (non-disjunctive) event $\hat{O}$ in the surface structure is "replaced" by a disjunctive event $\check{\mathbf{O}} = \{\hat{O}^n \mid n \leqslant N\}$ in the N-multiplied structure; the propensity of $\mu(e^* \rightarrowtail \hat{O})$ (in the surface structure) is identified with the propensity $\mu'(\langle e^*, 0\rangle \rightarrowtail \check{\mathbf{O}})$ (in the N-multiplied structure), the latter being equal to $\sum_{n \leqslant N} \mu'(\langle e^*, 0\rangle \rightarrowtail \hat{O}^n)$, by the properties of propensity functions. In accordance with the central tenet of our theory of propensities, these latter propensities should fully supervene on the propensities of appropriate sets of *causae causantes*. So, we should have: $(*)$ $\mu(e^* \rightarrowtail \hat{O}) = \sum_n \mu'(CC(\langle e^*, 0\rangle \rightarrowtail \hat{O})^n)$.

In the context of the Bell-Aspect experiment, the formula $(*)$ above requires one to calculate $\mu'(\langle e^*, 0\rangle \rightarrowtail (a_{i,m} \cup b_{j,m})^n)$ by using the larger set given by Eq. (8.17), whereas in our derivation of the Bell-CH inequality we used the smaller set specified by Eq. (8.18). That is, in the structure with a probabilistic hidden variable for the PFB in Bell-Aspect set-up, the contribution of the agent-based transitions, $a \rightarrowtail a_i$ and $b \rightarrowtail b_j^n$, is ignored. The discrepancy between the two sets of *causae causantes* shows a gap in our derivation: the condition of Adequate Propensity Assignment does not apply correctly in the context of the Bell-Aspect set-up, as it distorts the causal situation in question. Can one salvage the derivation despite this discrepancy? Although we only reluctantly assign propensities to agent-based transitions, we have to assign propensity to the large set of Eq. (8.17), which includes agent-based transitions, in order to understand what this discrepancy involves.

Let us thus assign a propensity μ' to the mentioned set. By Postulate 7.2, μ' is then defined for all sets obtained by varying i, j, m, m', and n in the mentioned formula. Let us focus on the last step of our derivation, "multiply by μ' and sum", and on Eq. (8.14). To justify this step , μ' needs to factor in a three-fold way:

$$\begin{aligned} p_{ij}^n((X_{i,ij}^n = m) \wedge (X_{j,ij}^n = m')) = \\ \mu'(\langle e^*, 0\rangle \rightarrowtail H^n,\ (a_i \rightarrowtail a_{i,m})^n,\ (b_j \rightarrowtail b_{j,m'})^n,\ (a \rightarrowtail a_i)^n,\ (b \rightarrowtail b_j)^n) = \\ \mu'(\langle e^*, 0\rangle \rightarrowtail H^n) \cdot \mu'((a_i \rightarrowtail a_{i,m})^n,\ (b_j \rightarrowtail b_{j,m'})^n) \cdot \mu'((a \rightarrowtail a_i)^n,\ (b \rightarrowtail b_j)^n), \end{aligned} \tag{8.19}$$

and the following propensities should be constant:

$$\begin{aligned}&\mu'((a \rightarrowtail a_i)^n,\ (b \rightarrowtail b_j)^n) = \mu'((a \rightarrowtail a_{i'})^n,\ (b \rightarrowtail b_{j'})^n) = \\ &= K(n),\ \text{with} \sum_{n \leqslant N} K(n) = 1.\end{aligned} \tag{8.20}$$

Significantly, note that the choice points in $\mathscr{W}'_{BA}$ are divided into C and E in such a way that the counterparts of each a and b belong in C, whereas $\langle e^*, 0\rangle$ and counterparts of each a_i and of each b_j are in E. The former serve as initials of agents-based transitions, whereas the latter are initials of Nature-given transitions. Given this division, C/E independence (see Def. 8.23) implies the following factorization, for every allowable i, j, m, m' and n:

$$\begin{aligned}&\mu'(\langle e^*,0\rangle \rightarrowtail H^n,\ (a_i \rightarrowtail a_{i,m})^n,\ (b_j \rightarrowtail b_{j,m'})^n,\ (a \rightarrowtail a_i)^n,\ (b \rightarrowtail b_j)^n) = \\ &\mu'((a \rightarrowtail a_i)^n,\ (b \rightarrowtail b_j)^n) \cdot \mu'(\langle e^*,0\rangle \rightarrowtail H^n,\ (a_i \rightarrowtail a_{i,m})^n,\ (b_j \rightarrowtail b_{j,m'})^n).\end{aligned} \tag{8.21}$$

We next note that our finitistic version of the Markov Principle permits a further factorization: The premises of the Markov Principle are satisfied as e^* is below each a_i and each b_j in $\mathscr{W}_{BA}$, these ordering relations carry over to $\mathscr{W}'_{BA}$, and the relevant finitistic assumptions hold as well.

$$\begin{aligned}&\mu'(\langle e^*,0\rangle \rightarrowtail H^n,\ (a_i \rightarrowtail a_{i,m})^n,\ (b_j \rightarrowtail b_{j,m'})^n) = \\ &\mu'(\langle e^*,0\rangle \rightarrowtail H^n) \cdot \mu'((a_i \rightarrowtail a_{i,m})^n,\ (b_j \rightarrowtail b_{j,m'})^n).\end{aligned} \tag{8.22}$$

Putting the two factorizations together, we see that C/E independence together with the Markov Principle justifies the three-fold factorization of Eq. (8.19), which in turn, taken together with the assumption stated in Eq. (8.20), justifies the last step of the derivation of the Bell-CH inequality (i.e., to multiply by μ' and to sum over). It is important to stress that we needed C/E independence to derive this inequality. Finally, now observe that the constancy of Eq. (8.20) implies Parameter Independence.[28]

[28] Given the causal discrepancy described above, the "big space" approach looks more attractive, since it assigns μ-propensity to sets like in Eq. (8.16), and consequently the condition of Adequate Propensity Assignment has a proper causal underpinning. Nevertheless, the problem with obtaining the three-fold factorization (Eq. 8.19) needed for the derivation of the Bell-CH inequality persists.

8.4.4.6 Consequences from our analysis

Let us reflect on the meaning of the above findings. The main message concerns the causal situation of the Bell-Aspect set-up. The set-up includes agent-based transitions, and this has consequences both for the sets of *causae causantes* and for the propensities. Given these agent-based transitions, there is a mismatch between propensities used in the analysis and the sets of *causae causantes* operating in the set-up. The mismatch can nevertheless be bridged by accepting some intuitive-looking postulates. In this context, the idea of C/E independence comes to the fore, which, given that propensities are assigned to agent-based transitions, receives the precise formulation of Def. 8.23. Given the Markov Principle (in our innocuous finitistic formulation), C/E independence justifies the use of settings-independent propensities $\mu'(\langle e^*, 0\rangle \rightarrowtail H^n)$ in the derivation of the Bell-CH inequality. This settings-independent measure captures the condition of No Conspiracy. There were two more assumptions, Outcome Independence assumed in the notion of a structure with a hidden variable for PFB, and the constancy of propensities of Eq. (8.20), which is tantamount to Parameter Independence. Our postulates, which govern propensities and often have a causal motivation, thus entail the standard premises of Bell's theorem (i.e., Outcome Independence, Parameter Independence, and No Conspiracy). To have a full list of our postulates, we used Outcome Independence, Eq. (8.20) (Parameter Independence), C/E propensity independence, and the Markov Principle. This set arguably gives us a better insight into what is implicated by the failure of the Bell-CH inequality than the standard set of premises, as there are some differences in status between our postulates. On the one hand, we have the high-level claims of C/E independence and the (finitistic) Markov Principle. They have a high-level status as no experiment-based argument for them looks feasible; on the other hand, it is hard, if not impossible, to conceive of a world without them. The high-level status of these two postulates suggests that they should be retained. In contrast, Eq. (8.20) looks more like a testable statement, as it concerns agents' propensities to choose alternative settings in the experiment. Agents can be trained to make unbiased choices, which arguably informs about propensities of the relevant agent-based transitions: they could be made numerically the same. We concede, however, that this observation is not very persuasive, as words like 'testing' and 'training' have limited sense in the deep realm of hidden variables.[29] In any case,

[29] A popular argument for accepting Parameter Independence (see Jarrett, 1984; Shimony, 1984), which relates to superluminal communication, can also be rehearsed in the BST framework. Its gist

given the empirical overtones of Eq. (8.20), and assuming that our everyday observations about agents carry over to their working in the realm of hidden variables, our recommendation is to accept it as well. The option that remains is to reject Outcome Independence, which *prima facie* agrees with the majority view. The novelty that BST brings is that a failure of Outcome Independence now concerns propensities. In non-technical language, that failure means that on a deep level, the Nature-given propensities fail to factor. That is, the degree of possibility of a complex happening does not supervene on the degrees of possibility of its components. In contrast, there is independence of agent-based transitions and Nature-given transitions, as expressed via C/E independence. This is the lesson about modalities that Bell's theorem brings, if it is analyzed from a BST perspective. We hope that with this lesson we are somewhat closer to achieving the great task that J. S. Bell (1997, p. 93) once posed.

> I think you must find a picture in which perfect correlations are natural, without implying determinism, because that leads you back to nonlocality. And also . . . as far as our individual experiences goes, our independence of the rest of the world is also natural. So the connections have to be very subtle . . .

8.5 Exercises to Chapter 8

Exercise 8.1. Prove Fact 8.4 (i.e., show that the generic-extended structure of a given BST_{92} surface structure is a non-empty, dense, strict partial).

Hint: Use Fact 8.3 and the definition of the ordering.

Exercise 8.2. Prove Theorem 8.1 (i.e., show that the generic-extended structure of a given BST_{92} surface structure is a BST_{92} structure).

Hint: Use Lemma 8.1 and the definition of the ordering.

Exercise 8.3. Prove Fact 8.11.

Hint: See Appendix B.8 for a proof based on the finiteness of S.

is that, if an observer knows how Eq. (8.20) fails, and knows "in which" hidden variable she is, she can learn in a faster than light way what settings her partner was choosing.

Exercise 8.4. Construct a BST_{92} surface structure $\langle W, <, e^*, E, C\rangle$, a non-contextual instruction set $\lambda \in \mathfrak{I}_n$, and a directed set $A \subseteq W$ with the following property: For every $a \in A$, there is $h_a \in H_a \subseteq \mathrm{Hist}(W)$ such that h_a matches λ, but there is no $h^* \in \mathrm{Hist}(W)$ such that $A \subseteq h$ and h^* matches λ.

Hint: Assume that A contains an infinite chain $\{a_i\}_{in \in \mathbb{N}}$ of binary choice points such that the chain occurs provided each a_i has outcome $+$. Let $E = \{e\}$, with e a binary choice point that is *SLR* to the chain, and $T = e \rightarrowtail +$. The CFB results from the assumption that any initial finite segment of pluses on the chain is consistent with T, but all the pluses on the chain are not consistent with T. There might be more choice points needed to secure that the construction is a BST structure.

Exercise 8.5. Show that the non-contextual instruction set λ of Eq. (8.2) on p. 259 is in fact maximal.

Exercise 8.6. Prove Theorem 8.5 (i.e., prove that the N-multiplied structure corresponding to a probabilistic BST_{92} surface structure $\mathscr{W}$ with a designated event e^*, as defined by Def. 8.18, is a BST_{92} structure as well).

Hint: Use Def. 8.18 of the ordering $<'$ and the form of histories in $\mathscr{W}'$ established via Lemma 8.4.

Exercise 8.7. Prove Lemma 8.5 (i.e., prove that a BST_{92} structure corresponding to a probabilistic BST_{92} surface structure is a *probabilistic* BST_{92} structure in the sense of Def. 7.3).

Hint: Note that any set $Y \subseteq \mathrm{TR}(W')$ on which μ' is defined either corresponds to a subset of $\mathrm{TR}(W)$ or has the form $\{\langle e^*, 0\rangle \rightarrowtail H^m\} \cup T^n$, where T^n corresponds to some $T \subseteq \mathrm{TR}(W)$. Assume that μ satisfies Postulates 7.1, 7.2, 7.3, and 7.4. Then use the conditions of an adequate propensity assignment of Def. 8.20 together with Def. 8.18 of the ordering $<'$ to show that μ' satisfies these Postulates as well.