



Cite this article: Krämer L, del Rio L. 2018
Operational locality in global theories. *Phil.
Trans. R. Soc. A* **376**: 20170321.
<http://dx.doi.org/10.1098/rsta.2017.0321>

Accepted: 16 April 2018

One contribution of 17 to a discussion meeting
issue 'Foundations of quantum mechanics and
their impact on contemporary society'.

Subject Areas:

mathematical physics

Keywords:

locality, causality, non-signalling

Author for correspondence:

Lidia del Rio

e-mail: delrio@phys.ethz.ch

Operational locality in global theories

Lea Krämer¹ and Lidia del Rio^{1,2}

¹Institute for Theoretical Physics, ETH Zurich, Switzerland

²School of Physics, University of Bristol, Bristol, UK

LdR, 0000-0002-2445-2701

Within a global physical theory, a notion of locality allows us to find and justify information-processing primitives, like non-signalling between distant agents. Here, we propose exploring the opposite direction: to take agents as the basic building blocks through which we test a physical theory, and recover operational notions of locality from signalling conditions. First, we introduce an operational model for the effective state spaces of individual agents, as well as the range of their actions. We then formulate natural secrecy conditions between agents and identify the aspects of locality relevant for signalling. We discuss the possibility of taking commutation of transformations as a primitive of physical theories, as well as applications to quantum theory and generalized probability frameworks. This 'it from bit' approach establishes an operational connection between local actions and local observations, and gives a global interpretation to concepts like discarding a subsystem or composing local functions.

This article is part of a discussion meeting issue 'Foundations of quantum mechanics and their impact on contemporary society'.

1. Introduction

In modelling local agents acting within a global theory, the intuitive assumption is that both their actions and their knowledge are restricted to a bounded region. The canonical example is a scientist who has full control of her laboratory and can perform local tomography. In reality though, the breadth of knowledge and the range of action of agents may be decoupled. For example, prisoners can acquire global knowledge by reading the news, but their actions are limited to small subsystems.

© 2018 The Authors. Published by the Royal Society under the terms of the Creative Commons Attribution License <http://creativecommons.org/licenses/by/4.0/>, which permits unrestricted use, provided the original author and source are credited.

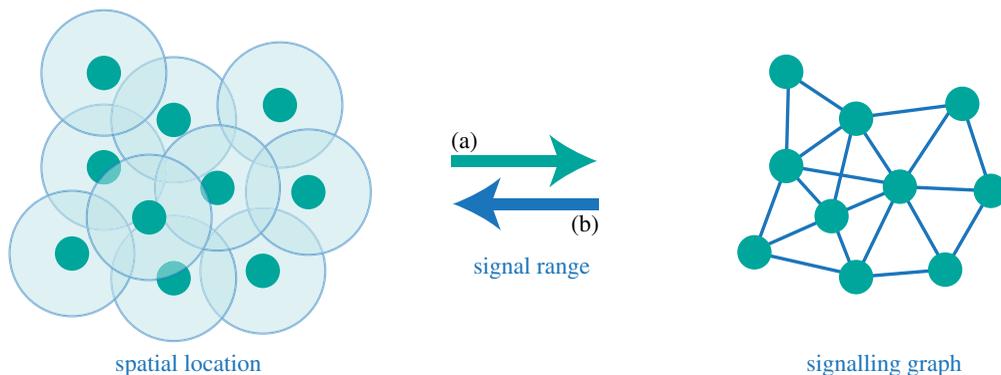


Figure 1. There and back again: physical locality and signalling. On the left, the spatial location of several agents (dark dots) and their range of communication (overlapping circles) are depicted; on the right, the corresponding signalling graph. (a) In physical theories, the notion of a space–time background where agents are positioned, together with principles about the range of signalling (e.g. the finite speed of light), allows us to derive information-processing concepts like non-signalling agents. (b) Reverse direction: starting from the notion of agents that may or may not be able to communicate and minimal assumptions on the nature and range of signalling, it may be possible to deduce both the space–time structure of the theory and the position of agents in it, to a good approximation. We can take inspiration from a simple example in the field of localization in wireless sensor networks (for a review, see, for example, [4]). Forest fire prevention mechanisms can be implemented by dropping a large number of smoke-detecting sensors from a plane over the forest. The sensors (our agents) are equipped with short-range communication systems, and land at random positions. One then collects the data of which sensors can signal to each other. From the signalling graph, it is possible to reconstruct the relative positions of the sensors on the ground to high accuracy—that way, when the smoke alarm goes off in a sensor, the fire-response team can quickly locate it. (Online version in colour.)

Conversely, someone locked in a control room may only have local knowledge of the shapes of different buttons, but pressing one may have global consequences. The observation that the knowledge and action do not always go hand in hand implies that in order to model agents we have to specify both (§2). This naturally leads us to search for minimal operational constraints needed to ensure that agents are truly local.

Here, we motivate a notion of *secrecy* between agents, which captures whether actions performed by an agent (like writing a message, choosing a bit or preparing a quantum state) can be perceived by another (§3); traditional notions of non-signalling correspond to an extended secrecy between space-like separated regions (§5). This work brings together and clarifies concepts of locality used in quantum theory, generalized probabilistic theories and field theories. It highlights that the *state space* and *transformations* of a theory are but a subjective choice of representation of the underlying physical theory from a viewpoint that is convenient to a given agent, as argued by Spekkens [1]. Here, we tentatively suggest commutation of transformations as a primitive of physical theories. In particular, we show how to derive local agents (and effective descriptions of local subsystems) from commutation relations on global transformations (§4).

This work draws from our ‘resource theories of knowledge’ [2], and has natural applications in multi-player settings, such as cryptographic scenarios, games or resource theories. There is yet a more exciting possible application: to recover the space–time structure of a physical theory from the primitive notion of test agents, in the spirit of Hardy’s operational general relativity [3] and of the task of localization in wireless sensor networks [4]. The idea is to send out agents (or probes) to unknown positions, see if they can communicate with each other, and use the signalling graph to define distances between agents, reconstruct their relative positions and infer properties of space–time (figure 1). For this, we must first find appropriate, theory-independent notions of agents and signalling.

2. Modelling agents

We start with a top-down approach, where we first describe a global theory (as seen by a global agent), and then model restricted agents acting within that theory.

(a) Global theory

From the point of view of a given global agent, a global theory may be represented via a state space Ω and a set of transformations \mathcal{T} that are available to the agent [5–8]. We can think of the state space as the ‘language’ chosen by this global observer to describe nature. For example, Ω could be the set of coordinates and momenta of all celestial bodies; in quantum theory, it could be the set of valid density matrices over a global Hilbert space. It need not be a static picture: in astronomy, an alternative state space Ω' could be the set of possible trajectories of celestial bodies, and in quantum theory it could include all global Hamiltonians that determine the free evolution of density matrices. Three observations are pertinent at this point: firstly, Ω is not the ultimate description of reality, just a convenient representation from the point of view of a global agent; secondly, different pictures, like Ω and Ω' , may be related and mapped to one another [1,2]; and thirdly, Ω need not have any special structure *a priori* besides being a set—indeed, the approach laid out here will allow us to find an operational subsystem structure in the set of states.

The transformations in \mathcal{T} represent all actions that the theory allows the global agent to implement. We can think of them as the ways in which the agent may test a theory, by applying actions that change state parameters. For example, an explicit theory of a quantum universe may allow only for unitary operations, while a more generous theory could equip the agent with implicit large ancillas, and allow her to implement general quantum channels, state preparations and even tomography. Again the two views can be related: the latter is an *effective theory* derived from the unitary quantum theory, by internalizing part of the global space as belonging to the agent and her instruments, and not to the object of study (the rest of the universe) [2]. In the context of field theories, this is discussed as *emerging agency* [3]. In a superdeterministic theory, there is only one possible course of evolution for the universe, and \mathcal{T} consists only of functions that apply it (for example, $\mathcal{T} \cong \{e^{-iHt}\}_t$ where the global agent is given some choice of time). Formally, \mathcal{T} is a monoid of functions $f: \Omega \rightarrow \Omega$: it contains the identity transformation and is closed under concatenation (an associative binary operation), such that performing two actions subsequently, $f \circ g$, is still an allowed operation. We discuss the monoidal assumption and possible relaxations in §5.

(b) Local agents

Local agents are characterized by limited knowledge: their inability to distinguish global states that appear identical in their eyes. We can formalize this by building equivalence classes of states that are indistinguishable from the perspective of an agent. For example, in quantum theory, we could have an agent Bob who only has access to a Hilbert space \mathcal{H}_B ; two global states are indistinguishable (or equivalent) from Bob’s perspective if they have the same marginal in \mathcal{H}_B . This defines an equivalence relation $\sigma \sim_B \rho : \text{Tr}_{\bar{B}}\sigma = \text{Tr}_{\bar{B}}\rho$, where $\text{Tr}_{\bar{B}}$ denotes a partial trace over all systems except B . The corresponding equivalence classes are

$$[\rho]_B := \{\sigma \in \Omega : \text{Tr}_{\bar{B}}\sigma = \rho_B\}.$$

Taking the quotient over, this equivalence relation gives us a new space state Ω/\sim_B , which is in one-to-one correspondence with the set of all reduced density matrices in \mathcal{H}_B . This is Bob’s *effective state space*, sufficient to encode all the information that he can observe about any global state (figure 2). In this case, the map from the global to the local spaces (the *canonical map*) is given by the partial trace

$$\mathbf{h}_B : \Omega \rightarrow \Omega/\sim_B,$$

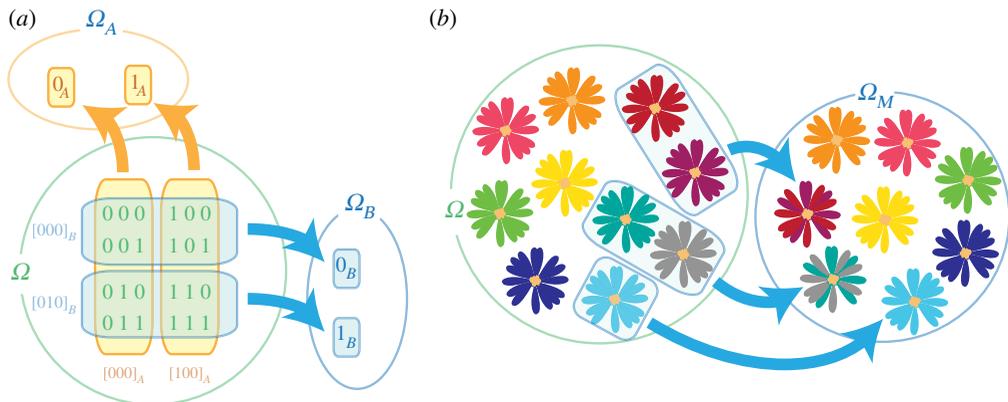


Figure 2. Building an agent's effective state space. The different states of a global space Ω are shown to an agent, who finds equivalence classes of (subjectively) indistinguishable states. Their effective state space is then the quotient space. (a) The global state space Ω consists of three bits, in the eight possible states depicted. An agent Alice can only see the first bit, therefore she cannot distinguish the states in each vertical box (her equivalence classes $[000]_A$ and $[100]_A$). Her effective state space $\Omega_A = \Omega / \sim_A$ has only two states, which can be relabelled as 0_A and 1_A for convenience. Another agent Bob identifies the equivalent classes $[000]_B$ and $[010]_B$, which leads us to conclude that he can only see the second bit. Note that, for example, if Alice were able to apply transformations that only change the first bit, she could not signal to Bob (because he could not detect the change). (b) Here, Ω is the space of colours, which were shown to a partly colourblind agent Marco. Marco identified the colours that he could not distinguish, which allowed us to build his reduced state space of colours Ω_M .

$$\rho \mapsto [\rho]_B \cong \text{Tr}_{\bar{B}} \rho = \rho_B.$$

More generally, we can always build the effective state space of an agent in this way, even if we do not know anything about the structure of the global space (for instance whether it can be split into a convenient tensor form $\mathcal{H}_A \otimes \mathcal{H}_B$). The construction of an agent's effective state space $\Omega_B := \Omega / \sim_B$ is in the spirit of the Leibniz principle of identity of indiscernibles [9]. Yet, this operational procedure emphasizes that both discernibility and identity are subjective concepts (figure 2). Limitations on Bob's perspective may have nothing to do with spatial locality. Bob might only have access to crude measurement instruments unable to distinguish microscopic details of states, or he may not be able to distinguish a global phase or gauge [3]. In generalized probability frameworks, Bob's perspective can correspond to a grouping of individual global outcomes into events (appendix E). In algebraic quantum field theory, these equivalence classes could emerge from algebras of local observables (e.g. [10] for a review).

The other ingredient needed to define an agent, as we saw in the Introduction, is a description of the actions available to him. As his actions may have a global impact, a minimal approach is to take them to be a submonoid $\mathcal{T}_B \subseteq \mathcal{T}$ of the globally allowed transformations. We discuss relaxations of this definition in §5. Generalizations of this approach can be found in [2]. There, we also study explicit ways to move between global and local views (technically, related by *Galois insertions*), effective theories and other properties of local agents.

Definition 2.1 (Global theory and restricted agents). A *global theory of agents* is defined by a pair (Ω, \mathcal{T}) , where Ω (the *state space*) is a set and \mathcal{T} is a monoid of transformations $f: \Omega \rightarrow \Omega$, with the concatenation operation \circ .

A *restricted agent B* acting within the theory is defined by a pair (\sim_B, \mathcal{T}_B) , where \sim_B is an equivalence relation in Ω and \mathcal{T}_B is a submonoid of \mathcal{T} called the set of *local operations* of the agent. The quotient space $\Omega_B := \Omega / \sim_B$ is called the *effective space* of agent B. The *reduction* to the effective space is given by the canonical map

$$\begin{aligned} \mathbf{h}_B: \Omega &\rightarrow \Omega_B, \\ \rho &\mapsto [\rho]_B. \end{aligned}$$

Having defined the effective state space of a restricted agent, we can also see how such an agent perceives the outcome of her local actions. Namely, the effect of an action $f_B \in \mathcal{T}_B$ applied to a global state ρ is seen by agent B as $[f_B(\rho)]_B$. This has been explored in more detail in [2], and is not used in this work. In short, such an agent could then define a local theory of accessible states and actions, $(\Omega_B, \tilde{\mathcal{T}}_B)$, where the transformations $\tilde{\mathcal{T}}_B$ act on the reduced state space Ω_B as

$$\begin{aligned} \tilde{f}_B: \Omega_B &\rightarrow \Omega_B, \\ [\rho]_B &\mapsto \tilde{f}_B([\rho]_B) := \bigcup_{\omega \in [\rho]_B} [f_B(\omega)]_B. \end{aligned}$$

We can always further coarse-grain the effective state space of a given agent B in order to obtain a more restricted agent C . For example, in renormalization group flow, lowering the cut-off corresponds to coarse-graining over more and more observables [11,12]. The following proposition formalizes this idea [2, prop. III.5]. All proofs can be found in appendix A.

Proposition 2.2 (Nested agents). *Let (Ω, \mathcal{T}) be a global theory, and B, C two restricted agents. Then the following are equivalent:*

- (i) C has more restricted knowledge than B , that is, $[\rho]_B \subseteq [\rho]_C, \forall \rho \in \Omega$,
- (ii) there exists an equivalence relation $\sim_{B \rightarrow C}$ in B 's effective state space Ω_B such that $\Omega_C \cong \Omega_B / \sim_{B \rightarrow C}$.

3. Secrecy between agents

(a) Secrecy

Having defined agents, we may study conditions for secrecy and non-signalling between them. Consider a set-up of two agents Alice and Bob, represented by $A = (\sim_A, \mathcal{T}_A)$ and $B = (\sim_B, \mathcal{T}_B)$. Imagine that Alice wants to keep her actions (like writing a message or preparing a state) secret from Bob. This is achieved if Bob cannot tell whether she applied them, even after post-processing.¹

Definition 3.1 (Secrecy). We say that an agent A has access to *secret operations* $\mathcal{T}_A^S \subseteq \mathcal{T}_A$ towards another agent B if

$$f_B \circ g_A(\rho) \sim_B f_B(\rho)$$

for all $\rho \in \Omega$, $g_A \in \mathcal{T}_A^S$, $f_B \in \mathcal{T}_B$. If all actions in \mathcal{T}_A are secret towards B and those in \mathcal{T}_B are secret towards A , we say that the two agents are *mutually secret*.

We may ask if this definition is robust enough, that is, whether further pre- or post-processing by Alice and Bob could destroy the secrecy of a choice of action $g_A \in \mathcal{T}_A^S$. The next proposition shows that no matter how many 'secret' transformations in \mathcal{T}_A^S Alice implements, or how Bob acts in between to try and recover information, he will not detect any of the effects of Alice's actions. In addition, it is easy to see that pre-processing with a global function (such as distributing entanglement between the two parties) cannot lift secrecy, as definition 3.1 requires it to hold for all initial states.

Proposition 3.2 (Robustness of secrecy). *If A has secret operations \mathcal{T}_A^S with respect to B (according to definition 3.1), then pre- and post-processing cannot lift the secrecy, that is,*

$$\begin{aligned} f_B^N \circ g_A^N \circ \cdots \circ f_B^2 \circ g_A^2 \circ f_B^1 \circ g_A^1 \circ f(\rho) \\ \sim_B f_B^N \circ \cdots \circ f_B^2 \circ f_B^1 \circ f(\rho), \end{aligned}$$

for all states $\rho \in \Omega$, secret operations $\{g_A^i\}_i \subseteq \mathcal{T}_A^S$ and $\{f_B^i\}_i \subseteq \mathcal{T}_B$, and global operations $f \in \mathcal{T}$ and $N \in \mathbb{N}$.

¹Bob's effective space may include his local processing (states that I can distinguish after applying all my accessible operations) or not (states that I distinguish immediately, before further processing). For the sake of generality, we leave the freedom in this decision up to the agent, and account for post-processing in the definition of secrecy.

(b) Extended secrecy

We may also ask whether Alice's actions stay secret to Bob in the presence of an additional global transformation $f \in \mathcal{T}$. Transformations, such as a subsystem swap or a communication channel, may break secrecy; others, like the use of a Popescu–Rohrlich box, do not.² For this situation, we define an extended notion of secrecy in the spirit of definition 3.1, which reduces to definition 3.1 in the case $f = \text{id}$. Here, Bob may try to post-process information before and after the global transformation.

Definition 3.3 (Extended secrecy). Let A be an agent with access to secret operations towards an agent B , $\mathcal{T}_A^S \subseteq \mathcal{T}_A$. We say that \mathcal{T}_A^S is in addition secret (towards B) in the presence of a global transformation $f \in \mathcal{T}$ if

$$f_B \circ f \circ f'_B \circ g_A(\rho) \sim_B f_B \circ f \circ f'_B(\rho),$$

for all $\rho \in \Omega$, $g_A \in \mathcal{T}_A^S$, $f_B, f'_B \in \mathcal{T}_B$. We say that the agents are *mutually secret in the presence of f* if all actions in \mathcal{T}_A are secret towards B in the presence of f and vice versa.

We can now show that, analogously to proposition A.2, further pre- and post-processing by Alice and Bob cannot lift the secrecy.

Proposition 3.4 (Robustness of extended secrecy). *If an agent A only uses secret operations $g_A \in \mathcal{T}_A^S$ with respect to agent B in the presence of $f \in \mathcal{T}$, then further pre- and post-processing cannot lift the secrecy, that is,*

$$\begin{aligned} & \left(\bigcirc_{i=1}^N f_B^i \circ g_A^i \right) \circ f \circ \left(\bigcirc_{i=1}^N f_B^i \circ g_A^i \right) \circ g(\rho) \\ & \sim_B \left(\bigcirc_{i=1}^N f_B^i \right) \circ f \circ \left(\bigcirc_{i=1}^N f_B^i \right) \circ g(\rho) \end{aligned}$$

for all states $\rho \in \Omega$, local operations $\{g_A^i\}_i \subseteq \mathcal{T}_A$ and $\{f_B^i\}_i \subseteq \mathcal{T}_B$, and global operations $g \in \mathcal{T}$ and $N \in \mathbb{N}$.

In particular, for the case in which Bob only implements post-processing at the very end, proposition 3.4 implies that \mathcal{T}_A^S forms a monoid.

Corollary 3.5 (Secret monoid). *The set \mathcal{T}_A^S of secret operations in the presence of a global function $f \in \mathcal{T}$ forms a monoid, i.e. $\text{id} \in \mathcal{T}_A^S$ and*

$$f_A, g_A \in \mathcal{T}_A^S \implies f_A \circ g_A \in \mathcal{T}_A^S.$$

Naturally, if we further restrict the actions and knowledge of one of the agents (as in proposition 2.2), secrecy is maintained.

Corollary 3.6 (Restricted agents and secrecy). *Let A, B and C be three agents, such that C is more restricted than B , that is, $\mathcal{T}_C \subseteq \mathcal{T}_B$ and $[\rho]_B \subseteq [\rho]_C$, for all $\rho \in \Omega$.*

If \mathcal{T}_B was secret towards A (in the presence of $f \in \mathcal{T}$), the same is true of \mathcal{T}_C . If \mathcal{T}_A was secret towards B (in the presence of f), it is still secret towards C (idem).

4. Commuting agents

Now, we explore how secrecy is affected when the actions of two agents A and B commute. This is particularly relevant in the context of the non-signalling principle, because actions at space-like separation naturally commute.

Definition 4.1 (Commuting agents). We say that two agents A and B commute if

$$f_B \circ g_A(\rho) = g_A \circ f_B(\rho),$$

for all $\rho \in \Omega$, $g_A \in \mathcal{T}_A$, $f_B \in \mathcal{T}_B$.

²In generalized probability theories, PR boxes can be seen as transformations that take classical inputs and return outputs (appendix E).

For example, in field theory commutativity holds for measurements or field interactions at space-like separation, and this is in general how causality is recovered there [13].³ Motivated by this, we here take the commutation of actions in space-like separated regions as a fundamental building block in deriving agents that are secret relative to each other. Note that, in particular, finding commuting sets of transformations in \mathcal{T} is something that can be done prior to definitions of local agents; this is shown explicitly in [2].⁴ Commutation is also an operational property of the theory: for example, commutation is independent of the choice of reference frames in relativity and quantum field theory [13,14]. If two agents commute, secrecy follows from simpler conditions.

Proposition 4.2 (Secrecy for commuting agents). *If A and B commute, then if there exists a subset of actions $\mathcal{T}_A^S \subseteq \mathcal{T}_A$ such that $\forall \rho \in \Omega$, $g_A \in \mathcal{T}_A^S, f_B \in \mathcal{T}_B$,*

$$f_B \circ f \circ g_A(\rho) \sim_B f_B \circ f(\rho),$$

then \mathcal{T}_A^S is secret towards B in the presence of f . In particular, $g_A(\rho) \sim_B \rho$ for all $g_A \in \mathcal{T}_A, \rho \in \Omega$ implies secrecy of A towards B .

(a) Secrecy from commutation

Starting only from commutation relations on the global transformations, we can construct descriptions of local agents that have secret actions with respect to each other. More specifically, given any two commuting submonoids $\mathcal{T}_A, \mathcal{T}_B \subseteq \mathcal{T}$, we can construct equivalence relations \sim_A, \sim_B so that two agents Alice (\sim_A, \mathcal{T}_A) and Bob (\sim_B, \mathcal{T}_B) have secret actions with respect to each other.

The first step is to start with transformations \mathcal{T}_A ('Alice's transformations'), and look for the most generous effective state space Ω_A that is insensitive to transformations in \mathcal{T}_A . This will model the perspective of an agent, Bob, who cannot detect Alice's actions.⁵

Definition 4.3 (Perspective insensitive to transformations). Let $\mathcal{T}_A \subseteq \mathcal{T}$ be a submonoid of transformations. First, we define a binary relation \sim'_A in Ω called *convergence through \mathcal{T}_A* as

$$\rho \sim'_A \sigma \iff \exists f_A, g_A \in \mathcal{T}_A \text{ s.t. } f_A(\rho) = g_A(\sigma).$$

We take the transitive closure \sim_A of \sim'_A to define the *perspective insensitive to transformations \mathcal{T}_A* ,

$$\begin{aligned} \rho \sim_A \sigma &\iff \exists n \in \mathbb{N}, \{\tau_i\}_{i=1}^n \subseteq \Omega : \\ \rho &\sim'_A \tau_1 \sim'_A \tau_2 \sim'_A \dots \sim'_A \tau_n \sim'_A \sigma. \end{aligned}$$

The above construction gives us minimal restrictions for independent agents. The following theorem is adapted from [2].

Theorem 4.4 (Deriving secret agents). *Commuting submonoids $\mathcal{T}_A, \mathcal{T}_B \subseteq \mathcal{T}$ give rise to descriptions of mutually secret agents*

$$A = (\sim_B, \mathcal{T}_A) \quad \text{and} \quad B = (\sim_A, \mathcal{T}_B).$$

Indeed, all agents whose actions commute with \mathcal{T}_A and for whom transformations in \mathcal{T}_A are secret must be described by a coarse-graining of \sim_A (figure 3). This and related minor results can be found in appendix D. In appendix F, we generalize theorem 4.4 to extended secrecy in the presence of global functions. There, we also extend the construction of the effective spaces of two agents to the case where the two monoids of transformations do not commute: without commutation, this construction is not as simple.

³The simplest illustration of this is the commutation of the Klein–Gordon field operators $\phi(x)$ and $\phi(y)$ at space-like separated x and y , $[\phi(x), \phi(y)] = 0$. Such a commutation condition is also referred to in field theory as the *locality postulate* [14].

⁴Commutation relations result in a nice algebraic structure—a lattice—in the space of transformations [2]. This is also the case for the von Neumann bicommutant in operator algebras [15].

⁵Essentially, this perspective identifies sets of global states that Alice can locally make 'converge' to the same state.

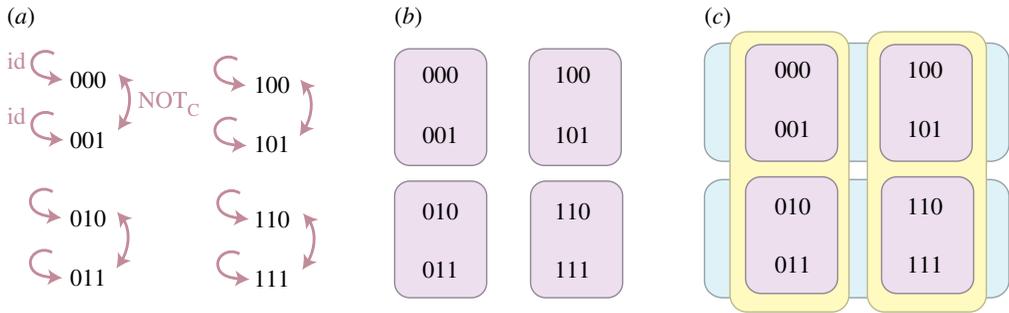


Figure 3. Three-bit example. Consider again the theory described by the state space Ω of three bits, and all transformations on those bits. (a) All operations \mathcal{T}_C that change the third bit (of which id and NOT_C are labelled). (b) Equivalence classes $[x]_f$ built according to definition 4.3. These correspond to the view of an agent who can only distinguish the first two bits. The equivalence relation \sim_f , which coarse-grains over the functions applied to the third bit, gives us the largest effective state space relative to which functions in \mathcal{T} are secret. (c) More coarse-grained equivalence classes $[x]_A$ (vertical, yellow) and $[x]_B$ (horizontal, blue), corresponding to an agent A who can only distinguish the first bit and an agent B who only sees the second bit, respectively. Operations in \mathcal{T}_C are still secret relative to these two agents. In addition, operations on the first bit are secret towards B and vice versa. These smaller effective state spaces correspond to equivalence relations on the effective state space Ω_f (as in the nested agents of proposition 2.2). The two-bit space Ω_f is a *common state space* of A and B , including states that could be distinguished if the two agents could work together, with $[x]_f = [x]_A \cap [x]_B$. (Online version in colour.)

(b) Perceived commutation from secrecy

We can now ask if the actions $\mathcal{T}_A, \mathcal{T}_B \subseteq \mathcal{T}$ of two mutually secret agents must always commute. The answer is no, not at a global level: unbeknown to the two agents, their actions could affect other degrees of freedom of the global theory. This can become relevant when the actions of two agents affect a common environment that is not directly accessible to them but could be recovered by a third party.

For example, consider again the state space of three bits, where Alice can only see the first bit and Bob the second. Now imagine that Alice has access to all the transformations that change the first bit and, as a side effect, reset the third bit to 0, while Bob has access to all the actions that act on the second bit and, as a side effect, flip the third bit. From a global viewpoint, their actions do not commute. However, for someone who only had access to the combined knowledge of Alice and Bob (the first two bits), their actions would appear to commute. For such an agent, only local time ordering of Alice and Bob's actions matters, as the two processes $f_A \circ f_B$ and $f_B \circ f_A$ are indistinguishable. This is yet another example of how subsystems and local descriptions represent simplified pictures of the global theory, reducing the degrees of freedom of the theory to an operational minimum for a given agent, who in this case would not need to model global time ordering.

5. Applications

In the previous sections, we have shown how to derive a notion of locality within a global theory starting from a primitive notion of individual agents, and their observed secrecy and commutation relations.

The operational approach laid out here has the advantage of carrying very few assumptions about the underlying physical theory. For example, it goes to a higher level of abstraction than generalized probability theories by not taking for granted that all agents express their knowledge in terms of reliable (classical) statistics about the outcomes of measurements.

Our notion of effective state spaces captures the concept of *beables* of a theory: aspects (or classes) of states that can in fact be physically observed and distinguished [3,16]. Our approach

highlights that beables are observer-dependent: for example, what appears to be a gauge may turn out to be only a local gauge [3], and the same applies to ‘global’ phases of quantum states or yet-to-be-discovered microscopic details of some structure. We can never rule out the existence of a more refined underlying theory, but with effective state spaces we can tailor the descriptions used in a theory to the level of detail needed for a particular application. This goes in the direction of the work of Colbeck & Renner [17], who showed that quantum theory is *complete* for the task of guessing measurement outcomes, and further refinements would be irrelevant.

As presented here, our framework simplifies the modelling of agents for the pedagogical purpose of highlighting the advantages of this general direction. In appendix C, we show how one could relax some of our assumptions to model agents that are limited in time or who can only approximately distinguish states. In the following, we discuss further applications and relation to other work.

(a) Non-signalling

One natural application of our extended notion of secrecy is the traditional non-signalling condition. To see this, imagine that the two agents are cooperating, so that Alice is trying to communicate information to Bob by means of some action $g_A \in \mathcal{T}_A$ on her side. Bob can now either directly apply post-processing $f_B \in \mathcal{T}_B$, or he can wait for some time to pass, as represented by a function $u_t \in \mathcal{T}$ that implements global time evolution over time t . If Alice and Bob are mutually secret in the presence of u_t , for all $t \leq T$, we conclude that they cannot signal to each other in this time window.

In appendix E, we show explicitly how our notion of extended secrecy implies traditional non-signalling in the framework of generalized probabilistic theories [18,19] for the case of two parties performing binary measurements (two inputs, two outputs on each side), where the state space consists of probability distributions over outcomes of possible measurements on physical systems. The generalization to finitely many inputs and outputs is straight forward.

(b) Reconstructing space–time

Building on the example above, if two agents cannot signal in the presence of u_t for $t \leq T$ and in addition can signal in the presence of u_t for $t > T$, this can be used to define a *distance* between the two agents, via $d \propto T$. The proportionality constant can be interpreted as the speed of signal propagation, for example the speed of light.

The challenge to obtaining a meaningful distance is two-fold: firstly, choosing a ‘natural’ family of transformations $\{u_t\}_t$ to represent time evolutions, and secondly, choosing a family of agents that do not conflate different types of coarse-grainings. For example, locality and macroscopicity each give rise to a natural notion of distance, relating to the space between agents and to precision of observation, respectively; the latter could be used to quantify chaos given a family of time evolutions.

More generally, we can try to use signalling between agents to infer properties of space–time of a given theory, as illustrated in figure 1. Some steps in this direction have been given, for example, in [3,20,21]. This would be of particular interest in the context of field theories [10,22,23]. We leave the generalization of the operational approach depicted in figure 1 to reconstruct position as future work.

(c) Relation to modular approaches

Our global approach complements modular, bottom-up constructions [24], like process theories based on symmetric monoidal categories [5–8,25]. For the purpose of comparison with our work, modular theories can be understood as theories of individual systems (or ‘objects’) and local actions (processes) on those systems, which allow for parallel and sequential composition of processes on different systems. Typically, they assume the following. (1) Processes with matching

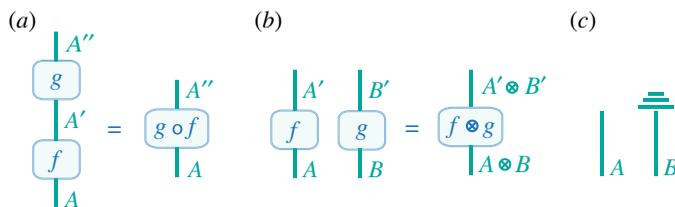


Figure 4. Process theories. Processes theories are modular, bottom-up constructions that can be faithfully represented by diagrams [5–8,25]. Lines represent *systems* (or ‘objects’) and boxes *processes* on systems: wires fed in from below a box can be understood as inputs to the process, while wires coming out of the top represent the outputs of the process. Diagrams can be composed due to the strong subsystem structure imposed on process theories, where actions are not seen as affecting the global space but explicitly associated with local systems. (a) Processes can be composed in sequence when their output and input systems match. (b) Processes can be composed in parallel on combined systems. (c) Discarding a subsystem (e.g. taking the partial trace) is indicated by three horizontal lines; in our approach this corresponds to coarse-graining over the relevant degrees of freedom (that is, going to a smaller effective space). Other conditions can be imposed: e.g. in [7], causal loops are forbidden, and outputs are always connected to inputs. (Online version in colour.)

output and input systems can be composed sequentially. That is, a process $f: A \rightarrow A'$ can be composed with a process $g: A' \rightarrow A''$, to form a new process $g \circ f: A \rightarrow A''$ satisfying

$$g \circ f(A) = g(f(A))$$

(figure 4a). (2) Any two systems A and B can be combined in parallel to form a composite system denoted by $A \otimes B$. (3) Any two processes $f: A \rightarrow A'$ and $g: B \rightarrow B'$ can be composed to yield a process $f \otimes g: A \otimes B \rightarrow A' \otimes B'$ satisfying

$$(f \otimes g)(A \otimes B) = f(A) \otimes g(B)$$

(figure 4b). This last assumption implies that processes act locally without disturbing other systems, and that actions on independent systems always commute. This allows us to represent process theories in terms of diagrams that can be easily composed (figure 4).

Our approach is more general in that we do not assume the strong subsystem structure imposed by conditions (2) and (3). As such, our work strengthens Coecke *et al.*'s [7] argument that non-signalling can be derived from a simpler condition (appendix B). In general, our top-down view can be taken as a precursor and sanity check for process theories. In complex global theories, a strong subsystem structure may not be clear-cut from the start. The cautious researcher can first use our approach to test different reduced descriptions for independence conditions. If she succeeds in finding independent effective spaces—which is not always possible—she may then frame them as subsystems and attempt a modular construction.

At a conceptual level, our approach gives a global interpretation to aspects of process theories that are more epistemic than physical. For example, if we think of subsystems as building blocks of a global space, it appears natural to see ‘discarding a subsystem’ as a physical action, like throwing away a piece of Lego (figure 4c). However, if we start from the global space and see subsystems as arbitrary restricted descriptions, then ‘discarding a subsystem’ corresponds to a coarse-graining over the relevant degrees of freedom (for example, going from Ω_{AB} to an effective space Ω_A), a change of perspective rather than a physical transformation.

(d) Relation to causal structures

Our notion of secrecy between agents is analogous to *causal independence* between events in graphs used to study causality in physics. Causal structures [1,26–30] try to capture the causal relations between events within a larger context (figure 5). Both causality (as expressed by Reichenbach’s principle) and secrecy are guiding principles of a certain way of representing a theory (causal structures and restricted agents, respectively) that help us understand a complex situation—they

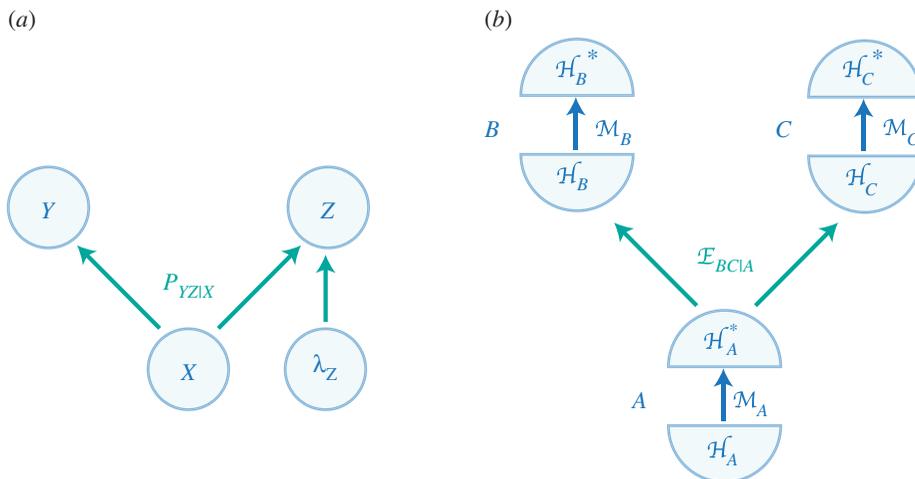


Figure 5. Causal structures. (a) In classical causal models, nodes are associated with random variables corresponding to events, while arrows carry causal influence, as specified by conditional probability distributions like $P_{YZ|X}$. These models and distributions may be extended if one later learns of additional causes, like λ_Z . (b) A model for quantum causal structures proposed in [31], where each node is associated with the intervention of an agent in a local space. A node i is represented by input and output Hilbert spaces, H_i and H_i^* , and by a quantum instrument \mathcal{M}_i that links the two and corresponds to the agent's intervention (for example, \mathcal{M}_i could be a local measurement followed by a local preparation depending on the outcome). Causal influence is explicitly carried by quantum maps like $\mathcal{E}_{BC|A}$, which acts like a channel from A to BC . (Online version in colour.)

are not necessarily fundamental features of the laws of Nature. How useful the representation is depends both on the guiding principle and on the choice of variables of interest (like events or agents).

Let us illustrate this. In classical causal graphs, events are represented by random variables, in principle subject to intervention (figure 5a). As we move from purely classical scenarios to more physical situations, like those involving quantum measurements, the formalism of causal structures is evolving to focus on agents and on explicit physical transformation as carriers of causal influence, similarly to our approach. For example, in the quantum causal structures of [31], events can correspond to quantum systems where agents can act locally (figure 5b). Generally speaking, 'events' embody a particular coarse-graining of a global picture into variables or subsystems of interest. As such, a single causal graph cannot reveal all the features of a complex theory—a different decomposition may explore new causal relations.⁶ The choice of relevant nodes can be guided by operationalism: (i) we start by picking 'variables' that we care about (like the outcomes of an experiment, or a subsystem corresponding to the perspective and range of intervention of an agent); (ii) we then use Reichenbach's principle and independence conditions to complete the causal graph, by identifying further nodes and constraining the channels between them.⁷ This procedure is similar in spirit to how in the present work we could start with the description of a few agents and use secrecy and commutation constraints to identify other subspaces and transformations of interest, or build a notion of locality. How successful we are in this endeavour depends largely on the (subjective) starting point—a poor initial choice of events or agents could make it impossible to find a meaningful causal graph or independent agents.

Even with a clever choice of initial variables, it could be that the guiding principle is not powerful enough to provide meaningful representations for all physical situations. This

⁶Furthermore, we can never know if we only have access to an effective state space to start with, and there is a deeper theory that changes all the causal relations, for example by providing new common causes.

⁷For example, in studying the process of coherent copy $\alpha|0\rangle_A + \beta|1\rangle_A \rightarrow \alpha|0\rangle_B|0\rangle_C + \beta|1\rangle_B|1\rangle_C$ in [31], we start with three nodes of interest (systems A , B and C), which are forced by Reichenbach's principle to complete the graph with a fourth node.

is likely to be the case in both approaches, which are still rooted in classical intuitions—resulting in concepts like agents, Reichenbach’s principle and time order. In trying to explain a physical scenario in terms of these classical notions, we risk running into paradoxes such as the inconsistencies between quantum agents in [32]. It remains to explore whether both our approach and causal models can handle this kind of physical challenge, and whether extensions to cover them would still be intuitive enough to help us make sense of the world.

Data accessibility. This article has no additional data.

Authors’ contributions. The authors contributed equally to this paper.

Competing interests. The authors declare that they have no competing interests.

Funding. L.K. acknowledges support from the European Research Council via grant no. 258932, the Swiss National Science Foundation through the National Centre of Competence in Research *Quantum Science and Technology* (QSIT), and the European Commission via the project *RAQUEL*. L.d.R. acknowledges support from ERC AdG NLST and EPSRC grant *DIQIP*, from the FQXi grant *Physics of the observer*, and from the Perimeter Institute for Theoretical Physics. Research at the Perimeter Institute is supported by the government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Economic Development & Innovation.

Acknowledgements. We thank Bob Coecke and Barry Sanders for discussions on discarding and non-signalling; Lucien Hardy, Ryszard Paweł Kostecki and Renato Renner for discussions on locality; John-Mark Allen and Mirjam Weilenmann for discussions on causality; David Jennings and Markus Müller for bringing up approximate distinguishability; Roger Colbeck for feedback on this manuscript; Miyuko for sanctuary; and Sandu Popescu for the mantra ‘one idea, one paper’. This one goes out to Matt and Rob—let us disagree again soon.

Appendix A. Proofs

Proposition A.1 (Nested agents). *Let (Ω, \mathcal{T}) be a global theory, and B, C two restricted agents. Then the following are equivalent:*

- (i) C has more restricted knowledge than B , that is $[\rho]_B \subseteq [\rho]_C, \forall \rho \in \Omega$,
- (ii) There exists an equivalence relation $\sim_{B \rightarrow C}$ in B ’s effective state space Ω_B such that $\Omega_C \cong \Omega_B / \sim_{B \rightarrow C}$.

Proof. For each direction:

1 \rightarrow 2. We build the equivalence relation in Ω_B as

$$[\rho]_B \sim_{B \rightarrow C} [\sigma]_B \iff [\rho']_C = [\sigma']_C, \\ \forall \rho' \in [\rho]_B, \sigma' \in [\sigma]_B.$$

As $[\rho]_B \subseteq [\rho]_C$ for all $\rho \in \Omega$, $\sim_{B \rightarrow C}$ is a well-defined equivalence relation, and the reduced space $\Omega_B / \sim_{B \rightarrow C}$ is in one-to-one correspondence with the space of the equivalence classes $[\rho]_C$.

2 \rightarrow 1. By assumption, the reduction \mathbf{h}_C is isomorphic to $\mathbf{h}_{B \rightarrow C} \circ \mathbf{h}_B$. Therefore, $[\rho]_C \cong [[\rho]_B]_{B \rightarrow C}$ and $[\rho]_B \subseteq [\rho]_C$. ■

Proposition A.2 (Robustness of secrecy). *If A has secret operations T_A^S with respect to B (according to definition 3.1), then pre- and post-processing cannot lift the secrecy, that is*

$$f_B^N \circ g_A^N \circ \dots \circ f_B^2 \circ g_A^2 \circ f_B^1 \circ g_A^1 \circ f(\rho) \\ \sim_B f_B^N \circ \dots \circ f_B^2 \circ f_B^1 \circ f(\rho),$$

for all states $\rho \in \Omega$, secret operations $\{g_A^i\}_i \subseteq T_A^S$ and $\{f_B^i\}_i \subseteq \mathcal{T}_B$, global operations $f \in \mathcal{T}$ and $N \in \mathbb{N}$.

Proof. We apply the secrecy condition (definition 3.1) multiple times. Define $\rho^{(j)} := (\bigcirc_{i=1}^j f_B^i \circ g_A^i) \circ f(\rho)$. Starting from the left-hand side, we have

$$\begin{aligned} (\bigcirc_{i=1}^N f_B^i \circ g_A^i) \circ f(\rho) &= f_B^N \circ g_A^N(\rho^{(N-1)}) \\ &\sim_B f_B^N(\rho^{(N-1)}) \\ &= f_B^N \circ f_B^{N-1} \circ g_A^{N-1}(\rho^{(N-2)}) \\ &\sim_B f_B^N \circ f_B^{N-1}(\rho^{(N-3)}) \\ &\vdots \\ &\sim_B f_B^N \circ \dots \circ f_B^1 \circ f(\rho). \end{aligned}$$

Proposition A.3 (Robustness of extended secrecy). *If an agent A only uses secret operations $g_A \in \mathcal{T}_A^S$ with respect to the agent B in the presence of $f \in \mathcal{T}$, then further pre- and post-processing cannot lift the secrecy, that is*

$$\begin{aligned} (\bigcirc_{i=1}^N f_B^i \circ g_A^i) \circ f \circ (\bigcirc_{i=1}^N f_B^i \circ g_A^i) \circ g(\rho) \\ \sim_B (\bigcirc_{i=1}^N f_B^i) \circ f \circ (\bigcirc_{i=1}^N f_B^i) \circ g(\rho), \end{aligned}$$

for all states $\rho \in \Omega$, local operations $\{g_A^i\}_i \subseteq \mathcal{T}_A$ and $\{f_B^i\}_i \subseteq \mathcal{T}_B$, global operations $g \in \mathcal{T}$ and $N \in \mathbb{N}$.

Proof. The proof is analogous to the proof of proposition A.2 and uses definition 3.3; we also define

$$\tilde{\rho} := f \circ (\bigcirc_{i=1}^N f_B^i \circ g_A^i) \circ g(\rho)$$

and

$$\rho^{(j)} := (\bigcirc_{i=1}^j f_B^i \circ g_A^i) \circ g(\rho).$$

Then

$$\begin{aligned} (\bigcirc_{i=1}^N f_B^i \circ g_A^i) \circ f \circ (\bigcirc_{i=1}^N f_B^i \circ g_A^i) \circ g(\rho) \\ \sim_B f_B^N \circ \dots \circ f_B^1(\tilde{\rho}) \quad [\text{secrecy}] \\ \sim_B f_B^N \circ \dots \circ f_B^1 \circ f \circ f_B^N(\rho^{(N-1)}) \quad [\text{Def. 3.3}] \\ \vdots \\ \sim_B f_B^N \circ \dots \circ f_B^1 \circ f \circ f_B^N \circ \dots \circ f_B^1 \circ g(\rho). \end{aligned}$$

Corollary A.4 (Restricted agents and secrecy). *Let A, B and C be three agents, such that C is more restricted than B, that is $\mathcal{T}_C \subseteq \mathcal{T}_B$ and $[\rho]_B \subseteq [\rho]_C$, for all $\rho \in \Omega$.*

If \mathcal{T}_B was secret towards A (in the presence of $f \in \mathcal{T}$), the same is true of \mathcal{T}_C . If \mathcal{T}_A was secret towards B (in the presence of f), it is still secret towards C (idem).

Proof. As $\mathcal{T}_C \subseteq \mathcal{T}_B$, it is secret towards A. This also restricts the post-processing that C can do, and since $\rho \sim_B \sigma \implies \rho \sim_C \sigma$, we have that \mathcal{T}_A is secret towards C. ■

Proposition A.5 (Secrecy for commuting agents). *If A and B commute, then if there exists a subset of actions $\mathcal{T}_A^S \subseteq \mathcal{T}_A$ such that, $\forall \rho \in \Omega$, $g_A \in \mathcal{T}_A^S, f_B \in \mathcal{T}_B$,*

$$f_B \circ f \circ g_A(\rho) \sim_B f_B \circ f(\rho),$$

then \mathcal{T}_A^S is secret towards B in the presence of f . In particular, $g_A(\rho) \sim_B \rho$ for all $g_A \in \mathcal{T}_A, \rho \in \Omega$ implies secrecy of A towards B.

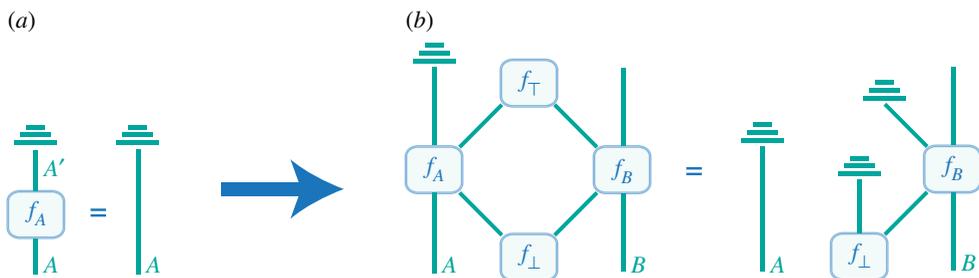


Figure 6. Terminality and non-signalling in process theories [7]. (a) Terminality: discarding a system A after applying a function f_A is the same as discarding the system directly. (b) Non-signalling: in this set-up, a global preparation process f_{\perp} distributes two systems to Alice and Bob, who perform local operations f_A and f_B , respectively, on these systems and additional inputs A and B . Given terminality, if system A is discarded after this, no information about the original input on A can travel to system B . This holds even in the presence of an effect f_{\top} acting on joint outputs of f_A and f_B . (Online version in colour.)

Proof. To show the second part of the proposition, secrecy of \mathcal{T}_A towards B , we have

$$\begin{aligned} f_B \circ g_A(\rho) &= g_A \circ \underbrace{f_B(\rho)}_{=:\rho' \in \Omega} \quad [\text{commutativity of } g_A, f_B] \\ &\sim_B f_B(\rho), \quad [\text{assumption: } g_A(\rho') \sim_B \rho'] \end{aligned}$$

for all $\rho \in \Omega, g_A \in \mathcal{T}_A, f_B \in \mathcal{T}_B$.

To show secrecy in the presence of f , we use

$$\begin{aligned} f_B \circ f \circ f'_B \circ g_A(\rho) &= f_B \circ f \circ g_A \circ \underbrace{f'_B(\rho)}_{=:\rho' \in \Omega} \quad [\text{commutativity of } g_A, f'_B] \\ &\sim_B f_B \circ f \circ f'_B(\rho), \quad [\text{assumption: } f_B \circ f \circ g_A(\rho') \sim_B f_B \circ f(\rho')] \end{aligned}$$

for all $\rho \in \Omega, g_A \in \mathcal{T}_A, f_B, f'_B \in \mathcal{T}_B$. ■

Lemma A.6. The perspective $\sim_{\mathcal{A}}$ induced by a submonoid $\mathcal{T}_A \subseteq \mathcal{T}$ is an equivalence relation in Ω .

Proof. By construction $\sim_{\mathcal{A}}$ is transitive, reflexive and symmetric. ■

Theorem A.7 (Deriving secret agents). Commuting submonoids $\mathcal{T}_A, \mathcal{T}_B \subseteq \mathcal{T}$ give rise to descriptions of mutually secret agents

$$A = (\sim_{\mathcal{B}}, \mathcal{T}_A) \quad \text{and} \quad B = (\sim_{\mathcal{A}}, \mathcal{T}_B).$$

Proof. We must show that

$$f_B \circ g_A(\rho) \sim_{\mathcal{A}} f_B(\rho),$$

for all $\rho \in \Omega, g_A \in \mathcal{T}_A$ and $f_B \in \mathcal{T}_B$. By proposition ??, we only need to show $g_A(\rho) \sim_{\mathcal{A}} \rho$, for all $\rho \in \Omega, g_A \in \mathcal{T}_A$. This holds since $\text{id} \in \mathcal{T}_A$ (as \mathcal{T}_A is a monoid), and so $g_A(\rho) \sim_{\mathcal{A}} \text{id}(\rho)$. We proceed analogously to find the effective state space of A . ■

Appendix B. Relation to terminality

In [7], Coecke argues that a process theory [8,25] is non-signalling if it satisfies a simpler condition dubbed *terminality*. Terminality states that local processes on a system right before ‘discarding’ it cannot have any observable effect (figure 6a). Discarding subsystems is a concept that corresponds to tracing out or coarse-graining over local information. For example, in quantum theory, it is implemented by the partial trace: terminality is naturally satisfied for completely positive trace-preserving operations on the discarded systems, but does not hold for non-deterministic effects such as projections onto particular outcomes of measurements [7].

In our language, the condition of terminality corresponds to the independence condition $f_A(\rho) \sim_B \rho$, where f_A are local functions on a system A and \sim_B corresponds to the local picture of other systems B outside A . Recall that the assumptions behind process theories like [7] impose some structure on transformations and agents, in particular, commutation between agents' local actions. As we saw in proposition 4.2, this independence condition together with commutation already implies secrecy.

It remains to see if our secrecy condition is equivalent to the non-signalling of [7], depicted in figure 6*b*. This non-signalling corresponds roughly to secrecy under pre- and post-processing, as implemented by an initial state preparation f_{\perp} and a deterministic effect f_{\top} . In our picture, robustness of secrecy under pre- and post-processing is ensured by proposition A.2. In this case, pre-processing with a function f_{\perp} can be included without loss of generality in the initial state. On the other hand, post-processing with f_{\top} is eliminated by the choice of local perspective \sim_B .

Hence, the result of [7] that terminality implies non-signalling follows from our propositions A.2 and A.5 together. Our premise that actions by different local agents commute is weaker than the assumptions employed by Coecke *et al.* [7]. In conclusion, our approach strengthens the argument in [7] for the significance of a condition like terminality. At the same time, we take a more general approach to subsystems than the bottom-up model of process theories in [7], thus highlighting the role of commutation in the context of defining local agents and non-signalling.

Appendix C. Relaxing some assumptions

Let us now give some guidelines on how to relax two of the assumptions of our framework, in order to cover more realistic representations of agents.

(a) Approximate distinguishability

Often agents may not have clear-cut distinguishability criteria. For example, an agent may categorize light frequencies into basic colours such as green and blue—there may be some frequencies that the agent could file as both green and blue. In the language of PBR [33], the reduced states 'blue' and 'green' would be epistemic and not ontic (with respect to the underlying state space of frequencies Ω). Agents could also have a notion of approximate distinguishability, for example, of the sort 'I can distinguish these two states with probability $1 - p$.'

We propose a simple approach to address these cases. (i) Build a generous effective state space Ω_B by assigning different reduced states to every two states in Ω that can be distinguished *in principle* by the agent. (ii) Build an *approximation structures* in the effective state space [2]. An approximation structure comprehends all neighbourhoods $\{\mathcal{B}^{\epsilon}(\rho)\}_{\rho \in \Omega_B, \epsilon \in \mathcal{E}}$ parametrized by whatever measure \mathcal{E} is operational for the agent. For example, one valid approximation structure for quantum states corresponds to the ϵ -balls induced by the trace distance; another could be just the *cover* {blue, green, ...} of the possible colours assigned to each frequency. (iii) Build notions of *approximate secrecy*, where we can demand, for example

$$\mathbf{h}_B \circ f_B \circ g_A(\rho) \in \mathcal{B}^{\epsilon}(\mathbf{h}_B \circ f_B(\rho)),$$

for all $\rho \in \Omega$, $g_A \in \mathcal{T}_A$, $f_B \in \mathcal{T}_B$, instead of the stricter condition of secrecy, where we demand that the two final states are completely indistinguishable from B 's perspective. The properties of approximate secrecy are inherited from the approximation structure.

(b) Time-limited agents

In this work, we model local actions as monoids \mathcal{T}_A and \mathcal{T}_B . When applying secrecy to find non-signalling conditions between time-limited agents, the monoidal structure of actions is only a convenient approximation, which allows us to concatenate post-processing actions indefinitely. The intuition behind this approximation is that Alice and Bob's actions can be implemented

essentially instantaneously, compared to the relevant time scales. One example would be the action of choosing a bit as an input to a measurement, by pressing a button in Alice's laboratory (see appendix E). With this interpretation, \mathcal{T}_A and \mathcal{T}_B can consistently be modelled as monoids, because it is assumed that the concatenation of two instantaneous actions can again be implemented instantaneously. In this model, time evolution is explicitly modelled by global functions $u_t \in \mathcal{T}$; this could include the actual effect of pressing the button.

When functions in \mathcal{T}_A and \mathcal{T}_B on Alice's and Bob's sides take some finite time $t > 0$ to implement, we may instead of full monoidal structure only have $f_A \circ g_A \in \mathcal{T}$ if the functions f_A and g_A together take less than a given time T to implement. In this case, the notion of secrecy or non-signalling and our results that relate to it can still be recovered for functions and concatenations of functions that do not exceed this time-frame T .

Appendix D. Commuting agents: additional results

In the main text, we have noted that the perspective \sim_A induced by transformations \mathcal{T}_A is minimal: any agent (\sim_B, \mathcal{T}_B) whose actions \mathcal{T}_B commute with \mathcal{T}_A and towards whom transformations in \mathcal{T}_A are secret must, in fact, be described by a coarse-graining of \sim_A .

Proposition D.8 (Induced perspective is minimal). *Let $B = (\sim_B, \mathcal{T}_B)$ be an agent towards whom \mathcal{T}_A is secret, and such that \mathcal{T}_A and \mathcal{T}_B commute. Then*

$$[\rho]_B \supseteq [\rho]_A, \quad \forall \rho \in \Omega,$$

with \sim_A the equivalence relation induced by \mathcal{T}_A . This implies that there exists an equivalence relation $\sim_{A \rightarrow B}$ in the effective state space Ω / \sim_A such that

$$\Omega_B \cong (\Omega / \sim_A) / \sim_{A \rightarrow B}.$$

Proof. As \mathcal{T}_A is secret towards B ,

$$\rho \sim_B g_A(\rho)$$

and so, due to transitivity of \sim_B , also

$$\exists f_A, g_A \in \mathcal{T}_A \quad \text{s.t. } f_A(\rho) = g_A(\sigma) \implies \rho \sim_B \sigma.$$

Again due to transitivity it directly follows that

$$\rho \sim_A \sigma \implies \rho \sim_B \sigma$$

and so

$$[\rho]_B \supseteq [\rho]_A,$$

for all $\rho \in \Omega$. We may thus employ proposition A.2. ■

Corollary D.9. *Let $\mathcal{T}_A, \mathcal{T}_B \subseteq \mathcal{T}$ be monoids such that $\mathcal{T}_A \subseteq \mathcal{T}_B$. Then the induced equivalence relations \sim_A and \sim_B satisfy*

$$[\rho]_B \supseteq [\rho]_A, \quad \forall \rho \in \Omega.$$

This again implies that there exists an equivalence relation $\sim_{A \rightarrow B}$ in the effective state space Ω / \sim_A such that

$$\Omega / \sim_B \cong (\Omega / \sim_A) / \sim_{A \rightarrow B}.$$

Proof. This follows from the fact that \mathcal{T}_A is secret towards \sim_B , together with proposition D.8. The second statement follows again from proposition A.2. ■

Finally, the following proposition shows that equivalence classes $[\rho]_A$ are preserved by commuting transformations \mathcal{T}_B , providing an operational interpretation to the perspective of agents (\sim_A, \mathcal{T}_B) : namely, states that are indistinguishable from Bob's point of view remain indistinguishable after he applies functions $f_B \in \mathcal{T}_B$,

$$\omega \sim_A \rho \implies f_B(\omega) \sim_A f_B(\rho).$$

Proposition D.10 (Induced perspective is operational). Let $\mathcal{T}_A, \mathcal{T}_B \subseteq \mathcal{T}$ be commuting transformations. Then the perspective \sim_A induced by \mathcal{T}_A satisfies

$$\omega \sim_A \rho \implies f_B(\omega) \sim_A f_B(\rho),$$

for any $f_B \in \mathcal{T}_B$ and $\omega, \rho \in \Omega$.

Proof. From the definition of \sim_A it follows that

$$\begin{aligned} \omega \sim_A \rho &\implies \exists n \in \mathbb{N}, \{\tau_i\}_{1 \leq i \leq n} \text{ with } \tau_i \in \Omega : \\ &\rho \sim'_A \tau_1 \sim'_A \tau_2 \sim'_A \dots \sim'_A \tau_n \sim'_A \omega \end{aligned}$$

with

$$v \sim'_A \omega \iff \exists f_A, g_A \in \mathcal{T}_A \text{ s.t. } f_A(v) = g_A(\omega)$$

for all $v, \omega \in \Omega$. But then, because \mathcal{T}_A and \mathcal{T}_B commute, for all of $\omega, \tau_1, \dots, \tau_n, \rho$ it holds that

$$\begin{aligned} v \sim'_A \omega &\iff \exists f_A, g_A \in \mathcal{T}_A \text{ s.t. } f_A(v) = g_A(\omega) \\ &\implies \exists f_A, g_A \in \mathcal{T}_A \text{ s.t. } f_B \circ f_A(v) = f_B \circ g_A(\omega) \\ &\iff \exists f_A, g_A \in \mathcal{T}_A \text{ s.t. } f_A \circ f_B(v) = g_A \circ f_B(\omega) \\ &\implies f_B(v) \sim'_A f_B(\omega) \end{aligned}$$

for all $f_B \in \mathcal{T}_B$. From the definition of \sim_A as the transitive closure of \sim'_A , it then also follows that

$$\omega \sim_A \rho \implies f_B(\omega) \sim_A f_B(\rho).$$

■

Appendix E. Application to generalized probability theories

Generalized probability theories (GPTs [18,19,34–36]) are a framework to infer as much as possible about a physical system without making assumptions about its inner workings (like the assumption that states can be represented as vectors in a Hilbert space). Instead, it is assumed that agents can implement and label a number of physical procedures, like preparations, transformations and, crucially, measurements. Note that the agents need not know the actual physical state prepared; in order to label a procedure, they only need to be confident that they can repeat it. Indeed, the basic assumption behind GPT frameworks is that agents can extract significant measurement statistics (for example, by repeating a procedure many times). Hence, GPTs model outputs of measurements as random variables, and agents' knowledge of procedures as probability distributions.

The usual approach to build GPTs is bottom-up, starting with local procedures that can be composed to reach a global theory. Here, we are interested in the opposite direction: given a global GPT, can we find meaningful notions of local variables? Firstly, we need to model global and local knowledge.

(a) Basic formalism

While we are inspired by known GPT models [18,19,34–36], we take a slightly different and simplified approach here. The idea is that agents only have direct access to classical random variables (like input settings and outputs of a physical measurement). As they correspond to accessible information, we denote probability distributions over these random variables by *states*.

Transformations f are naturally modelled by conditional probability distributions $P_{Z|X}^f$ that take

input to output states, such that $f(P_X) = P'_Z$, with

$$P'_Z(Z) = \sum_{x \in X} P^f_{Z|X}(z|x) P_X(x).$$

For example, suppose that we want to model an experiment where an agent performs a quantum measurement by pressing two buttons: button X prepares a quantum state ρ^x and button Y measures it according to the POVM $\{E^y_z\}_z$ with possible outcomes $\{z\}_{z \in Z}$. The distributions P_{XY} over inputs and P'_Z over outputs correspond to accessible 'states.' We model the transformation as a conditional distribution $P^f_{Z|XY}$ with $P^f_{Z|XY}(z|x, y) = \text{Tr}(E^y_z \rho_x)$. The final distribution P'_Z of outcomes given an input distribution P_{XY} is therefore

$$\begin{aligned} P'_Z(z) &= \sum_{x \in X} \sum_{y \in Y} P^f_{Z|XY}(z|x, y) P_{XY}(x, y) \\ &= \sum_{x \in X} \sum_{y \in Y} \text{Tr}(E^y_z \rho_x) P_{XY}(x, y). \end{aligned}$$

Note that it is the conditional distribution that encodes the 'physical' information about a particular setting (like the quantum state and POVM), which may be inaccessible to the agents.

To compare with the models of [18,19], our transformations are analogous to their states. However, in our agent-driven approach, we restrict the set of allowed measurements to those accessible to a particular agent in the resource theory—in this case, they can thus be seen as a subset of the *fiducial measurements* that define a state in [18,19].

In our model, *global states* correspond to distributions over a global random variable X . Restricted agents are those unable to distinguish some of the outcomes of the global variable. We can model this via arbitrary groupings of outcomes $x \in X$ into equivalence classes, i.e. *events* $\{B_b\}_b$. The *reduction* function \mathbf{h}_B to the effective state space of an agent B simply sums over all the probabilities of the individual outcomes $x \in B_b$ in each event B_b and returns the probability associated with the event,

$$P_B = \mathbf{h}_B(P_X), \quad \text{with } P_B(b) = \sum_{x \in B_b} P_X(x).$$

(b) Secrecy and non-signalling

Consider now two agents A and B whose actions commute. To guarantee secrecy of A towards B , we only need to satisfy the independence condition $g_A(P_X) \sim_B P_X$ (for all global P_X and all $g_A \in \mathcal{T}_A$, see proposition 4.2). In our language, this condition reads

$$\sum_{y \in B_b} \sum_{x \in X} P^{g_A}_{Y|X}(y|x) P_X(x) = \sum_{x \in B_b} P_X(x), \quad \forall b.$$

For simplicity, we took Y and X to be identical random variables that represent the global state before and after the transformation, and \mathbf{h}_B is a particular coarse-graining of outcomes into events. The condition then states that \mathbf{h}_B is insensitive to the transformation (conditional probability distribution) $P^{g_A}_{Y|X}$ from inputs $x \in X$ to outcomes $y \in Y$.

If A and B commute and are mutually secret, we can ask if an additional global transformation $f \in \mathcal{T}$ allows for signalling between them (definition 3.3). Our condition for extended secrecy in the presence of f , $f_B \circ f \circ g_A(P_X) \sim_B f_B \circ f(P_X)$, for all global P_X , $f_B \in \mathcal{T}_B$ and $g_A \in \mathcal{T}_A$, becomes

$$\begin{aligned} &\sum_{v \in B_b} \sum_{x, y, z} P^f_{V|Z}(v|z) P^f_{Z|Y}(z|y) P^{g_A}_{Y|X}(y|x) P_X(x) \\ &= \sum_{v \in B_b} \sum_{x, z} P^f_{V|Z}(v|z) P^f_{Z|X}(z|x) P_X(x), \quad \forall b. \end{aligned} \tag{E1}$$

Non-signalling functions are then those that do not let information encoded in $P^A_{Y|X}$ propagate to Bob's perspective \sim_B , that is, such that \mathcal{T}_A is secret with respect to (\sim_B, \mathcal{T}_B) in the presence of f .

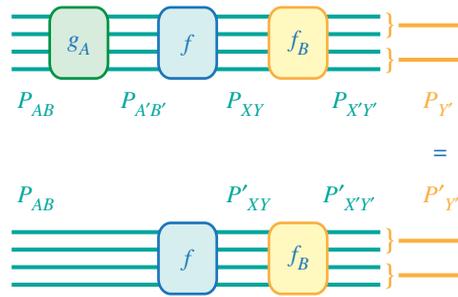


Figure 7. Extended secrecy in GPTs. Consider two agents Alice and Bob and a two-bit random variable AB (the four possible outcomes are represented by vertical lines). Our global state space consists of distributions over two bits (whose names we change after each step to make a proof more readable). Let the reduction to the state space of Bob be a coarse-graining over the first bit, $P_Y(y) = P_{XY}(0, y) + P_{XY}(1, y)$, as shown on top. Secrecy in the presence of a function f (equation (E 1)) corresponds to $P_{Y'} = P'_{Y'}$ in the diagram. This is satisfied, for example, if f represents the use of a PR-box, while g_A and f_B correspond to choices of inputs and post-processing: Bob cannot guess Alice's choice of input after the use of a PR-box. This is equivalent to the traditional notion of non-signalling [17] applied to the PR-box. (Online version in colour.)

Here, again for simplicity X, Y, Z, V were chosen as identical random variables, and $P_{Z|X}^f = P_{Z|Y}^f$ both represent the conditional probability distribution corresponding to f . For a simple example in a two-bit space (figure 7).

We can compare our definition of secrecy in the presence of f to traditional notions of non-signalling. Consider again the simple case of a two-bit input and output space of figure 7. The definition of non-signalling found, for example, in [17] reads $P_{Y|AB} = P_{Y|B}$, or

$$\sum_{x=0,1} P_{XY|AB}(xy|a=0, b) = \sum_{x=0,1} P_{XY|AB}(xy|a=1, b), \quad \forall b, y \in \{0, 1\}. \quad (\text{E } 2)$$

This condition formalizes the idea that Bob cannot learn anything about Alice's input a by looking solely at his output y and input b . In our framework, Alice's choice of input is encoded in a local transformation $g_A \in \mathcal{T}_A$ (for example, g_A^0 could correspond to pressing a button to choose input 0 and g_A^1 to choose 1), and therefore 'Bob's ignorance about Alice's input' translates to 'Bob's ignorance about Alice's action g_A '.

In the following, we establish a direct equivalence between these two notions of non-signalling in this simple case; we expect this equivalence to hold in more general settings. Let us first flesh out the assumptions behind the equivalence. A gentle warning: we have labelled all the intermediate bits differently 'to avoid confusion' (figure 7). As we assume *a priori* that Alice and Bob have mutual secrecy (without f), we take that g_A only acts locally on Alice's bit,

$$P_{A'B'|AB}^{g_A}(a', b' | a, b) = P_{A'|A}^{g_A}(a' | a) \delta(b', b),$$

so that

$$g_A(Q_{AB}(a, b)) = \sum_{a,b} P_{A'|A}^{g_A}(a' | a) \delta(b', b) Q_{AB}(a, b) = \sum_a P_{A'|A}^{g_A}(a' | a) Q_{AB}(a, b).$$

Similarly, Bob's post-processing is encoded in $f_B \in \mathcal{T}_B$ which we also assume to be truly local, that is

$$P_{X'Y'|XY}^{f_B}(x', y' | x, y) = P_{Y'|Y}^{f_B}(y' | y) \delta(x', x).$$

The final distribution $P_{Y'}$ for Bob becomes

$$\begin{aligned}
 P_{Y'}(y') &= \mathbf{h}_B \circ f_B \circ f \circ g_A(Q_{AB}(a, b)) \\
 &= \sum_a \mathbf{h}_B \circ f_B \circ f(P_{A'|A}^{g_A}(a'|a) Q_{AB}(a, b)) \\
 &= \sum_a \sum_{a',b} \mathbf{h}_B \circ f_B(P_{XY|A'B'}^f(x, y|a', b) P_{A'|A}^{g_A}(a'|a) Q_{AB}(a, b)) \\
 &= \sum_a \sum_{a',b} \sum_y \mathbf{h}_B(P_{Y'|Y}^{f_B}(y'|y) P_{XY|A'B'}^f(x, y|a', b) P_{A'|A}^{g_A}(a'|a) Q_{AB}(a, b)) \\
 &= \sum_a \sum_{a',b} \sum_y \sum_x P_{Y'|Y}^{f_B}(y'|y) P_{XY|A'B'}^f(x, y|a', b) P_{A'|A}^{g_A}(a'|a) Q_{AB}(a, b).
 \end{aligned}$$

The condition for secrecy in the presence of f , equation (E 1), is then

$$\begin{aligned}
 &\sum_{a,a',b,x,y} P_{Y'|Y}^{f_B}(y'|y) P_{XY|A'B'}^f(x, y|a', b) P_{A'|A}^{g_A}(a'|a) Q_{AB}(a, b) \\
 &= \sum_{a,b,x,y} P_{Y'|Y}^{f_B}(y'|y) P_{XY|AB}^f(x, y|a, b) Q_{AB}(a, b), \\
 &\forall y' \in \{0, 1\}, g_A \in \mathcal{T}_A, f_B \in \mathcal{T}_B, Q_{AB} \in \Omega.
 \end{aligned} \tag{E3}$$

Proposition E.11 (Equivalence to non-signalling in GPTs). *In the setting of figure 7, ‘our’ condition of non-signalling, equation (E 3), is equivalent to the ‘traditional’ notion, equation (E 2).*

Proof. In our language, the non-signalling condition of equation (E 2) reads

$$P_{XY|A'B'}^f(xy|b, a' = 0) = P_{XY|A'B'}^f(xy|b, a' = 1) =: P_{XY|A'B}^f(xy|b).$$

To show that equation (E 3) implies the above, we choose the particular local actions

$$P_{A'|A}^{g_A^0} = \delta(a', 0), \quad P_{A'|A}^{g_A^1} = \delta(a', 1) \quad \text{and} \quad P_{Y'|Y}^{f_B}(y'|y) = \delta(y', y).$$

There, g_A^0 corresponds to Alice’s choice of input 0, g_A^1 to her choice of 1, and f_B to no post-processing by Bob. These choices directly imply for all Q_{AB} ,

$$\sum_x P_{XY|AB}^f(xy|b, a = 0) Q_B(b) = \sum_x P_{XY|AB}^f(xy|b, a = 1) Q_B(b)$$

and so traditional non-signalling follows. For the other direction, we have simply

$$\begin{aligned}
 P_{Y'}(y') &= \sum_{a,a',b,x,y \in \{0,1\}} P_{Y'|Y}^{f_B}(y'|y) P_{XY|A'B'}^f(xy|a'b) P_{A'|A}^{g_A}(a'|a) Q_{AB}(ab) \\
 \text{non-signalling (equation (E 2))} &= \sum_{a,b,x,y} P_{Y'|Y}^{f_B}(y'|y) P_{XY|B}^f(xy|b) \underbrace{[P_{A'|A}^A(0|a) + P_{A'|A}^A(1|a)]}_{=1} \\
 &\quad \times Q_{AB}(a, b) \\
 &= \sum_{a,b,x,y} P_{Y'|Y}^{f_B}(y'|y) P_{XY|AB}^f(xy|b) Q_{AB}(a, b).
 \end{aligned}$$

■

This shows that, for example, PR boxes satisfy our definition of non-signalling functions. Examples for functions that are signalling are conditional probability distributions that swap the states on the two systems, or bitwise addition of the inputs a and b on the two sides into the outputs x and y .

Appendix F. Deriving secrecy without commutation

In principle, the way we have constructed an induced perspective \sim_A from a monoid \mathcal{T}_A can be extended to construct equivalence relations that yield secret agents (\sim_A, \mathcal{T}_A) and (\sim_B, \mathcal{T}_B) in the presence of a global function f , and even in the case where functions \mathcal{T}_A and \mathcal{T}_B do not commute. This is done in the following proposition, which, as is shown in the subsequent corollary, reduces to the definition of induced perspectives \sim_A and \sim_B in the case $f = \text{id}$ and commuting $\mathcal{T}_A, \mathcal{T}_B$.

Proposition F.12 (Deriving secret agents). *Let (Ω, \mathcal{T}) be a global theory, $\mathcal{T}_A^S, \mathcal{T}_B \subseteq \mathcal{T}$ be two monoids of transformations, and let $f \in \mathcal{T}$. Then the smallest equivalence class \sim_B on Ω towards which \mathcal{T}_A^S is secret in the presence of f ,*

$$f_B \circ f \circ f'_B \circ g_A(V) \sim_B f_B \circ f \circ f'_B(V),$$

for all $V \in S^\Omega$, $g_A \in \mathcal{T}_A^S, f_B \in \mathcal{T}_B$, is built as follows:

Define the relation \sim on Ω as

$$\rho \sim \sigma \iff \exists f_A, g_A \in \mathcal{T}_A^S \quad \text{s.t. } f_A(\rho) = g_A(\sigma).$$

Then define another relation \sim' as

$$\rho \sim' \sigma \iff \begin{cases} \rho \sim \sigma \quad \text{or} \\ \exists \rho', \sigma' \in \Omega, f_B \in \mathcal{T}_B \quad \text{s.t. } \rho = f_B(\rho'), \sigma = f_B(\sigma'), \rho' \sim \sigma' \quad \text{or} \\ \exists \rho', \sigma' \in \Omega, f_B, f'_B \in \mathcal{T}_B \quad \text{s.t. } \rho = f_B \circ f \circ f'_B(\rho'), \\ \sigma = f_B \circ f \circ f'_B(\sigma'), \rho' \sim \sigma'. \end{cases}$$

Finally, the relation \sim_B on Ω is the transitive closure of \sim' , namely through

$$\rho \sim_B \sigma \iff \exists n \in \mathbb{N}, \tau_1, \dots, \tau_n \in \Omega \quad \text{s.t. } \rho \sim' \tau_1, \tau_1 \sim' \tau_2, \dots, \tau_n \sim' \sigma.$$

Proof. Both the relation \sim and \sim' are by construction reflexive and symmetric. The relation \sim_B is then by construction also transitive, and thus constitutes an equivalence relation. The relation \sim_B furthermore gives rise to the smallest perspective towards which \mathcal{T}_B is secret in the presence of f : $\omega = g_A(\rho) \implies \rho \sim_B \omega$. By construction then also $f_B \circ f \circ f'_B \circ g_A(\omega) \sim_B f_B \circ f \circ f'_B(\omega)$ and $f_B \circ g_A(\omega) \sim_B f_B(\omega)$, for all $\omega \in \Omega, f_B, f'_B \in \mathcal{T}_B, g_A \in \mathcal{T}_A^S$. ■

Corollary F.13. *In the case of commuting $\mathcal{T}_A^S, \mathcal{T}_B \subseteq \mathcal{T}$, the equivalence relation \sim_B constructed in proposition F.12 that gives rise to secrecy in the presence of f simplifies accordingly and can be constructed as follows.*

Define the relation \sim on Ω as

$$\rho \sim \sigma \iff \exists f_A, g_A \in \mathcal{T}_A^S \quad \text{s.t. } f_A(\rho) = g_A(\sigma).$$

Then define another relation \sim' as

$$\rho \sim' \sigma \iff \begin{cases} \rho \sim \sigma \quad \text{or} \\ \exists \rho', \sigma' \in \Omega, f_B \in \mathcal{T}_B \quad \text{s.t. } \rho = f_B \circ f(\rho'), \sigma = f_B \circ f(\sigma'), \rho' \sim \sigma'. \end{cases}$$

Then, the relation \sim_B on Ω is the transitive closure of \sim' , namely through

$$\rho \sim_B \sigma \iff \exists n \in \mathbb{N}, \tau_1, \dots, \tau_n \in \Omega \quad \text{s.t. } \rho \sim' \tau_1, \tau_1 \sim' \tau_2, \dots, \tau_n \sim' \sigma.$$

If in addition $f = \mathbb{I}$, the relation \sim_B simplifies to

$$\rho \sim_B \sigma \iff \exists n \in \mathbb{N}, \tau_1, \dots, \tau_n \in \Omega \quad \text{s.t. } \rho \sim \tau_1, \tau_1 \sim \tau_2, \dots, \tau_n \sim \sigma$$

with \sim as above. This recovers the construction of induced perspectives \sim_A in definition 4.3.

Proof. In the case when functions in \mathcal{T}_A^S and \mathcal{T}_B commute, we can see that

$$\begin{aligned}\rho \sim \sigma &\iff \exists f_A, g_A \in \mathcal{T}_A^S \text{ s.t. } f_A(\rho) = g_A(\sigma) \\ &\implies \exists f_A, g_A \in \mathcal{T}_A^S \text{ s.t. } f_B \circ f_A(\rho) = f_B \circ g_A(\sigma) \\ &\iff \exists f_A, g_A \in \mathcal{T}_A^S \text{ s.t. } f_A \circ f_B(\rho) = g_A \circ f_B(\sigma) \\ &\iff f_B(\rho) \sim f_B(\sigma),\end{aligned}$$

for all $f_B \in \mathcal{T}_B$. This implies the respective simplifications of the relation \sim_B , and recovers the induced perspective \sim_A of \mathcal{T}_A^S in the case of $f = \text{id}$. ■

References

- Spekkens RW. 2012 The paradigm of kinematics and dynamics must yield to causal structure. (<http://arxiv.org/abs/1209.0023>)
- del Rio L, Krämer L, Renner R. 2015 Resource theories of knowledge. (<http://arxiv.org/abs/1511.08818>)
- Hardy L. 2016 Operational general relativity: possibilistic, probabilistic, and quantum. (<http://arxiv.org/abs/1608.06940>)
- Cheng L, Wu C, Zhang Y, Wu H, Li M, Maple C. 2012 A survey of localization in wireless sensor network. *Int. J. Distrib. Sens. Netw.* **2012**, 1. (doi:10.1155/2012/962523)
- Coecke B, Fritz T, Spekkens RW. 2014 A mathematical theory of resources. (<http://arxiv.org/abs/1409.5531>)
- Fritz T. 2015 The mathematical structure of theories of resource convertibility I. (<http://arxiv.org/abs/1504.03661>)
- Coecke B, Hasuo I, Panangaden P. 2014 Terminality implies non-signalling. *Electron. Proc. Theor. Comput. Sci.* **172**, 27–35. (doi:10.4204/EPTCS.172.3)
- Coecke B. 2010 A universe of processes and some of its guises. In *Deep beauty: understanding the quantum world through mathematical innovation*, pp. 128–186. Electronic Proceedings in Theoretical Computer Science. (<http://arxiv.org/abs/1009.3786>)
- Leibniz G 1739 *Tentamina theodicææ de bonitate dei: libertate hominis et origine mali*. Frankfurt, Germany: C.H. Bergerus.
- Valente G. 2013 Local disentanglement in relativistic quantum field theory. *Stud. Hist. Philos. Sci. B: Stud. Hist. Philos. Modern Phys.* **44**, 424–432. (doi:10.1016/j.shpsb.2013.09.001)
- Wilson KG, Kogut J. 1974 The renormalization group and the ϵ expansion. *Phys. Rep.* **12**, 75–199. (doi:10.1016/0370-1573(74)90023-4)
- Polchinski J. 1984 Renormalization and effective Lagrangians. *Nucl. Phys. B* **231**, 269–295. (doi:10.1016/0550-3213(84)90287-6)
- Peskin ME, Schroeder DV. 1995 *An introduction to quantum field theory*. Reading, MA: Addison-Wesley.
- Banks T. 2008 *Modern quantum field theory: a concise introduction*. Cambridge, UK: Cambridge University Press.
- von Neumann J. 1930 Zur Algebra der Funktionaloperationen und Theorie der normalen Operatoren. *Mathematische Annalen* **102**, 370–427. (doi:10.1007/BF01782352)
- Bell JS. 2004 *Speakable and unspeakable in quantum mechanics—collected papers on quantum philosophy*, 2nd edn. Cambridge, UK: Cambridge University Press.
- Colbeck R, Renner R. 2016 The completeness of quantum theory for predicting measurement outcomes. In *Quantum theory: informational foundations and foils* (eds G Chiribella, RW Spekkens), pp. 497–528. Dordrecht, The Netherlands: Springer.
- Hardy L. 2001 Quantum theory from five reasonable axioms. (<http://arxiv.org/abs/0101012>)
- Barrett J. 2007 Information processing in generalized probabilistic theories. *Phys. Rev. A* **75**, 032304. (doi:10.1103/PhysRevA.75.032304)
- Hoehn PA, Mueller MP. 2014 An operational approach to spacetime symmetries: Lorentz transformations from quantum communication. (<http://arxiv.org/abs/1412.8462>)
- Cao C, Carroll SM, Michalakis S. 2016 Space from Hilbert space: recovering geometry from bulk entanglement. (<http://arxiv.org/abs/1606.08444>)

22. Halvorson H, Mueger M. 2006 Algebraic quantum field theory. (<http://arxiv.org/abs/0602036>)
23. Wolters SAM, Halvorson H. 2013 Independence conditions for nets of local algebras as sheaf conditions. (<http://arxiv.org/abs/1309.5639>)
24. Hardy L. 2012 The operator tensor formulation of quantum theory. *Phil. Trans. R. Soc. A: Math. Phys. Eng. Sci.* **370**, 20110326. L. Hardy. (doi:10.1098/rsta.2011.0326)
25. Coecke B. 2006 Kindergarten quantum mechanics: lecture notes. In *AIP Conference Proceedings*, pp. 81–98. New York, NY: AIP. See <https://www.cs.ox.ac.uk/people/bob.coecke/VaxjoProc.pdf>.
26. Penrose R. 1972 *Techniques in differential topology in relativity*. Society for Industrial and Applied Mathematics.
27. Pearl J. 2000 *Causality: models, reasoning, and inference*. Cambridge, UK: Cambridge University Press.
28. Masanes L, Müller MP, Augusiak R, Pérez-García D. 2013 Existence of an information unit as a postulate of quantum theory. *Proc. Natl Acad. Sci. USA* **110**, 16 373–16 377. (doi:10.1073/pnas.1304884110)
29. Ried K, Agnew M, Vermeyden L, Janzing D, Spekkens RW, Resch KJ. 2015 A quantum advantage for inferring causal structure. *Nat. Phys.* **11**, 414–420. (doi:10.1038/nphys3266)
30. Chiribella G, Spekkens RW. 2016 *Quantum theory: informational foundations and foils*. Dordrecht, The Netherlands: Springer.
31. Allen J-MA, Barrett J, Horsman DC, Lee CM, Spekkens RW. 2016 Quantum common causes and quantum causal models. (<http://arxiv.org/abs/1609.09487>)
32. Frauchiger D, Renner R. 2016 Single-world interpretations of quantum theory cannot be self-consistent. (<http://arxiv.org/abs/1604.07422>)
33. Pusey MF, Barrett J, Rudolph T. 2012 On the reality of the quantum state. *Nat. Phys.* **8**, 475–478. (doi:10.1038/nphys2309)
34. Wootters WK. 1986 Quantum mechanics without probability amplitudes. *Found. Phys.* **16**, 391–405. (doi:10.1007/BF01882696)
35. Mana PGL. 2003 Why can states and measurement outcomes be represented as vectors? (<http://arxiv.org/abs/quant-ph/0305117>)
36. Mana PGL. 2004 Probability tables. (<http://arxiv.org/abs/quant-ph/0403084>)