

Promise Theory and the Alignment of Context, Processes, Types, and Transforms

Jan A. Bergstra

janaldertb@gmail.com

Minstroom Research BV, Utrecht, The Netherlands

Mark Burgess

mark.burgess.oslo.mb@gmail.com

ChiTek-i AS, Oslo, Norway

Submitted: 26 June 2020

Revised: 25 February 2021

Abstract

Promise Theory concerns the ‘alignment’, i.e. the degree of functional compatibility and the ‘scaling’ properties of process outcomes in agent-based models, with causality and intentional semantics. It serves as an umbrella for other theories of interaction, from physics to socio-economics, integrating dynamical and semantic concerns into a single framework. It derives its measures from sets, and can therefore incorporate a wide range of descriptive techniques, giving additional structure with predictive constraints. We review some structural details of Promise Theory, applied to Promises of the First Kind, to assist in the comparison of Promise Theory with other forms of physical and mathematical modelling, including Category Theory and Dynamical Systems. We explain how Promise Theory is distinct from other kinds of model, but has a natural structural similarity to statistical mechanics and quantum theory, albeit with different goals; it respects and clarifies the bounds of locality, while incorporating non-local communication. We derive the relationship between promises and morphisms to the extent that this would be a useful comparison.

1 Introduction

Promise Theory, first outlined in 2005 [1], is a compositional approach to describing systems of interacting agents. It has had some practical successes, over

the past decade, thanks to its straightforward applicability to technology [2–19]. Like many approaches to modelling, for machines, processes, functions, and operations, it describes systems in terms of agents whose autonomous behaviours are composed from the bottom up. Being generic in form, and technology-independent, it has arguably been easier to carry over into different disciplines than more abstract mathematical ideas like Category Theory, where the barrier to entry is much higher; Promise Theory has also been applied, with some predictive-analytical capability, to socio-economic fields across a multitude of scales [20–23].

Promise Theory separates the dynamics and semantics of an interaction into three phases, none of which may be considered completely deterministic:

- The declaration of agent discovery and ‘intent’ (to be defined).
- The process of implementing that intent (internally).
- The assessment of the outcome by an observer (internally or externally).

It is a simple generalization of an approach inspired by classical and quantum field theories in physics (where one has charge, field or messenger carriers, and response, which may be local or non-local with respect to the point of origin). However, it emphasizes richer semantics and the different interpretations of scale, such as one finds in modern ‘cloud’ computing and socio-economic systems.

Alternative descriptions of interactions take a variety of differing points of view, which makes comparing and contrasting them non-trivial. All approaches have to limit their scope in some way. Category theory, which deals with data types and their relationships on many levels, has enjoyed a growing popularity since its proposal and inception in [24], but remains very difficult for the casual reader, and attempts to straddle an awkward divide between precision and applicability. Other ideas like Petri Nets and Process Algebras are also widely used, in certain fields, both as machine models [25] and as mnemonics for other equational systems [26]. The attention to structural completeness comes often at the expense of brittle rigidity. Finally, there is physics, with the longest standing battery of theories for ‘stuff that happens’. Promise Theory is still somewhat simple-minded mathematically, yet surprisingly revealing in terms of generic and elementary principles. It peddles a more physical picture of interactions, rather than describing the abstract typology of its data.

With these notes, we attempt to lay a simple bridge between these disparate languages, by showing how promises can be represented in a language of types, and states. We do not expect to convince sceptics of Promise Theory’s usefulness, but at least help to clarify what it is and is not, for the curious who are more steeped in other disciplines.

1.1 Promise and imposition models

In Promise Theory, the central objects are *agents*, which are the sources and receivers of what one calls ‘promises’ and ‘impositions’. They are abstract

‘places’ where necessary resources and processes for promise-keeping are localized. Agents can make promises and impose on one another to advertise their autonomous behaviours, attempt to induce cooperation, and adapt to other agents. For some purposes, this simple starting point is enough to go quite far in reasoning about systems. The result is not a conventional logic – indeed, it defies the traditional modal-logic formulations of force and necessity, and works more in the manner of a local semi-deterministic, but causal observability, giving it some natural similarities with Quantum Mechanics as a theory of *incomplete information*.

To delve beyond this shallow (albeit surprisingly constructive) level, more structure is needed. We can infer, for instance, that agents have to remember state involved in the keeping of promises, which involves memory and either encapsulating or partaking in *processes*, interior or exterior to them. These processes need not be simple Markov processes of first order [27], as is often assumed of dynamical models. Indeed, the detailing of process may not need to be explained at all as long as one can enumerate outcomes. This promises great flexibility.

We note in passing that The Actor Model [28] fits into the set of models that are trivially represented by Promise Theory, but is too close to programming to be a general modelling framework. We do not want to limit the form of message interaction to state-machine transition-inducing events, nor say anything too specific about the nature of agents, as this might be quite different on different scales and in different circumstances.

In this work, we shall discuss the most elementary kind of ‘technical’ promise, called a *Promise of the First Kind* [4]. This kind of promise may be used to express more complicated kinds of promise, such as all those used in common parlance, which we do not discuss here – so there is no intrinsic limitation in making this choice. It is helpful to sketch a brief rapid-fire overview of the Promise Theory ‘manifesto’ in preparation:

- The world is made up of agents, which partake in the keeping of ‘promises’. Agents can also attempt to impose on one another through ‘impositions’.
- Agents may be composed from other agents, by collaborative interactions, at all scales, but it is assumed that there is a possibly elementary ground state at which agents cannot or need not be subdivided further.
- Agents can assist one another, thereby composing themselves into clusters, by making promises about their own behaviours. This implies a directed semi-coupling between agents, which is drawn as labelled arrows between them. Promise arrows may, themselves, be composed in parallel and in series, possibly subject to scaling transformations of the agents.
- If necessary, agent processes may be characterized in terms of a set of interior states q_A , for agent A , which can undergo transition events $q_A \mapsto q'_A$, and such events also compose to form processes. Transitions (which are rarely drawn to avoid a confusion of arrows) may be spontaneous or stimulated by other events.

- A particular assignment of values to states may be called a *configuration*, which, in a particular encoding may be written as a function $Q(q_A)$ onto unspecified value types, or in a special (often preferred) representation to a state occupancy vector of real numbers, denoted by $\psi(q_A)$.

In passing, we note that the usual preference for real number quantitative ‘truth’ values, through a ψ -representation, relates to the widespread preference for logical or probabilistic interpretations, such as those used in Quantum Mechanics. This is related to the statistics of observation and the semantics of its interpretation. In computer engineering, a representation $Q(q)$ in terms of structured data is more common (a database schema).

- The possible paths a process can take through its ‘configuration space’ $\psi(q_A)$ define ‘trajectories’ formed from increments of the form

$$\Delta\psi : \psi \mapsto \psi'$$

- The terminating states of a process, whose goal it is to keep a promise, that is the *possible end states*, are called the *possible outcomes* (and sometimes the *desired end states*) for the promise.
- Agent ‘intentions’ thus become bundles or selections (subsets) of potential pathways from the current state to one or more of its possible outcomes.
- Agents make promises about their intentions to other agents by signalling them with messages.
- An agent keeps a promise by undertaking some interior process, local to itself. An agent can only keep its own promises, so a promise-keeping process is always on the interior of an agent that makes the promise.
- The only non-local process is that which leads to *message passing*. The channel over which a promise is communicated might not be the same as the channel over which the promise-keeping process communicates with agents. For example, electronic components make their promises via component catalogues ($50\Omega \pm 5\%$ or $20 \mu F$, etc.), but the processes which keep them are inside materials that form electric circuits.

The outcome of this manifesto results in a hierarchy of labelled directed graphs, which seems straightforward enough – yet it would be remiss not to mention how radically different a story it is from conventional computer science, and also classical physics, which more commonly deal with transitional arrows of two particular kinds:

- *Ballistic trajectories*, in which an arrow represents a vector or an imposed transition, triggered by arrivals of some token and carrying certain properties, such as in a classical momentum-carrying collisions or fluid flow vectors.

- *Mappings*, in which one describes the structure of the domains and co-domains involved in transformations – but which do not directly imply causation.

Promises do not fall into either of these types, but can form a representation of either if necessary.

1.2 A discrete theory of agent alignment

Promises venture out of a particular *cul de sac* which dogs mathematical logic – namely how completeness and deterministic precision, represented as types, may be in conflict with the representation of realism and scale. Physics is successful because it deals in approximations. Logics tend to over-constrain models to the point of being inflexible or even misleading, and unwieldy in their adherence to algebras, even with so-called ‘bounded rationality’. Promise Theory admits weaker forms of coupling – akin to fuzzy logics.

We need a way of classifying agents and their behaviours in ways that may be exactly or partially similar. The similarity of agents is subject to observability constraints (which is the meaning of relativity in science, no only of the Einsteinian type, but certainly including that). The bodies of promises, which label and perhaps explicate their asserted behaviours, may be compared by various mathematical tools to define a degree of similarity, that is what we shall call alignment of intent.

Promise Theory employs arrow diagrams to represent promise interactions, but not identically to the usage in other model frameworks. Arrows may often be how theories describe an order of causation with respect to a particular process. Some caution is always needed to interpret arrows. In many descriptions of the world, arrows are assumed to be forces, point responses, fields or channels of influence that are *embedded* within larger ‘spaces’ (for example vector, topological, etc). This may imply an overreach of assumption in such models that goes beyond what could be considered parsimonious. Embedding spaces are employed as scaffolding, say for coordinatizing phenomena, and are typically given an independent and implicit significance which is unwarranted from a parsimonious system perspective – one might say, a case of muddling representation with phenomenon. In physics, for instance, one begins with differential equations of motion, expressed with real valued coordinates (\vec{x}, t) in a representation that presumes the existence of a smooth differentiable manifold, but could equally be formulated in a different way, without that assumption and perhaps with fewer side-effects. This is not the case in Promise Theory – agents are not considered to be embedded in a space or subject to influences from without (that are not represented by their promises). Promise Theory is a bottom-up construction, rather than a top-down summarization by rules.

Using a language of promises, one can try to decouple a representation that generates dynamics from the specific implementation ‘technology’ it employs – or at least draw attention to the assumptions embodied by a choice. One typically starts with less specific promises in order to avoid representation until it

becomes essential. When a given representation has more freedoms than it needs to describe a process, it is under-constrained, and a number of workarounds are needed to patch this: for example the supplementary specification of ‘boundary conditions’ (advanced or retarded), ‘selection rules’ and identification of ‘forbidden transitions’, ‘random probabilistic events’, etc. The simple idea of Promise Theory is that these various approaches can be rationalized under a single common narrative about causal bindings of autonomous local processes.

2 Agents

Let A be an agent, from a collection of such. Particular agents A, S, R, \dots etc, thus belong to a multiset of objects of a bare agent type. Agents have no assumed structure, though if we want to formalize promises further, a minimum amount can be inferred. If agents make no promises at all, then they are *a priori* indistinguishable. By labelling them differently, we assume that they have distinguishing promises. A name is thus a shorthand for a promise to all agents:

$$A_i \equiv A \xrightarrow{+\text{name}=i} *.$$
 (1)

Note, we may or may not assume that all agents can read the label! Colour blind agents might not be able to distinguish certain names, for instance. Agents are therefore indistinguishable until they identify themselves by making some kind of promise, which is accepted by another and then acts as a label or name tag that other agents can discriminate. In physics, for instance, we label the agents (usually called particles) with properties that can be observed and used by other agents. This is how force models work. Such labels may be considered to be assumed promises of the first kind, so that a promise view exposes that assumption explicitly.

2.1 Advertising process roles and properties

Classes of agent promise similar behaviours. We define the equivalence class to be a *role*, expressing dynamical or semantics similarity – like commodity components for sale in a market. However, the full nature of agents is both undefined and unknown until explained through the revelation of promises. Agents are, in fact, only observable (and thus identifiable and distinguishable) by the promises they make. Compatibility with other theories of computation and information implies some basic necessities about the internal structures of agents (Figure 1).

An agent A must have internal resources that include memory to keep its promises. This is a significant departure from a particle view to a process view of elementary structure – a shift from local space with global time to fully local spacetime. Memory is required to represent states q_A , which include the declaration of promises π_A made by A , and interior changes required to keep the promise (since an agent may only promise its own state outcomes). Interior processes further imply the existence of an ‘internal clock’ T_A , which ticks at

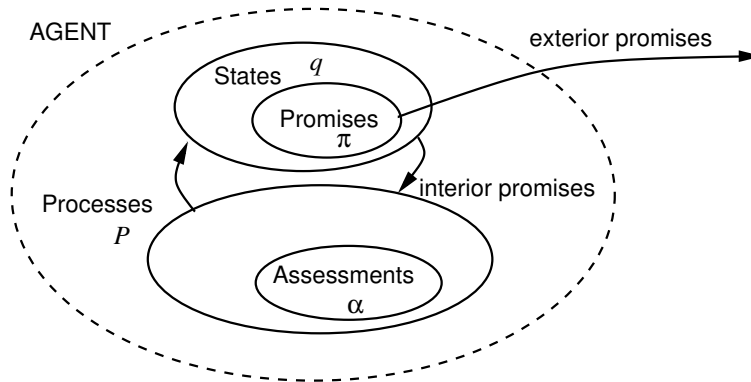


Figure 1: The structure of agents in Promise Theory requires that agents have interior (memory) states and processes by which promises are kept. The processes include ‘methods’ or ‘algorithms’ for keeping promises and for assessing outcomes of local promises and bindings to other agents. Processes must also be involved in making selections from alternatives that is ‘intentional behaviour’.

its own rate in agent A , since a process is a sequence of distinguishable changes that may form the basis for defining a clock. The processes required for minimal Promise Theory semantics include *assessments* $\alpha_A(\pi)$, or processes which assess whether promises have been kept, locally or remotely.

2.2 A scalable notion of trust

A generalized notion of *trust* enters unavoidably at this point as one of the assessments agents can make about one another [29]. Trust is inevitably aligned with the concept of promises in a human sense, because the extent to which an agent R might be willing or unwilling to accept a promise from S is affected by its trust in the promiser agent, with respect to the particular kind of promise [29]. The generalization of trust to more elementary mechanisms and scales is less mysterious than one might presume; it serves as a form of reliability assessment, which one can explain either statistically based on evidence (Frequentist), or as a boundary condition on belief (Bayesian). Trust, however, is a slightly mysterious quantity in practice. We think of it as is a potentially *ad hoc* assessment by an agent R , which is related to the belief that an agent S is likely to keep a promise π_S , and is one of most significant dependent assessments an agent A can make about another agent A' , since agents may decide to make, keep, withdraw, or respond to promises on the basis of a simple updated level of trust. Thus, the frequency view of probability falls away because it is deterministic and we can make no such assumption.

At the microscopic level, we do not normally use the term trust between agents, because we think of trust as a human assessment; but this need not be the case. In short, it is a parameter that influences the strength of coupling for

promise assessments. It plays a role similar to a coupling constant in physics (see the discussion in [5]).

In other words, between people, trust should not be thought of as a rationally assessable quantity, but something more like an empirical sampling frequency or probability that an agent will keep its promise, but rather as a classifier of ‘belief’ in whether an agent’s promise is sound, in the absence of validating information. Later evidence may revise this assessment up or down, depending on the agent’s assessment process. Quantitative modellers will inevitably want to map this assessment to some real number for comparison (as we have discussed in [29]), but what Promise Theory emphasizes is that this is not a global or unique prescription; rather, it is an *ad hoc choice by each individual agent*.

3 Boundaries

The definition of an agent implies the introduction of a kind of spacetime ‘boundary’ (see the dotted line in Figure 1) which delimits the ‘interior’ from ‘exterior’ of the agent. Interior may be logical or physical (indeed, if such a distinction can be considered). Boundaries are essentially discontinuous changes in where promises are in scope. The characteristics of interior and exterior dynamics may or may not share characteristics. Moreover, as agents compose to form ‘superagents’, forming a larger effective boundary from smaller ones, what was exterior on one scale might become interior on another. Notice how scales are promise type specific in a Promise Theory, a phenomenon we refer to as a semantic scale.




Figure 2: An agent A signals an exterior promise to another agent A' . Each agent has interior states q_A of undisclosed nature. Superagents may also contain sub-agents, whose promises on the interior of the superagent, may acts as some of the states of the superagent, and which, in turn, have their own interior states, forming a hierarchy. Agents are distinguishable by other agents only by the assessment of promises they make or keep.

Assuming a ‘ground state’ for the most atomic level of agents, that is that they are not infinitely divisible, then each boundary defines a notion of agent scale by the level of composition [30]. Composing agents of scale S leads to agents of scale $S + 1$, for any agent collection presumed elementary. At each

new scale, new promises can arise that may be considered independent of the promises they are composed from.

3.1 Agent scale, interior and exterior

On a large enough scale ($S \gg 1$) one could not reasonably expect to enumerate the many maps and transitions that have to take place, in order to trace the causal flow of every detail of agent interactions, nor to predict the new levels of language that emerge in this scaling [30]. A coarse grained prescription is needed [30]. There is therefore a natural connection between the ‘complexity’ of agent behaviours and their scale. Each new scale can combine new promises and agents, forming new processes with possibly new alphabets for interaction. This implies that there are new phenomena in composition, so that a system is not simply the direct sum of its agents. For example, a promise about agents’ collective interior states being ‘all on the left’ or ‘all on the right’ has no meaning for electrons, but would apply to configurations in offices or football pitches. Similarly, properties of agents like axial symmetry, cephalization, head to tail order, etc has well-defined meaning at scale but not at the elementary level.

The issue of boundaries further plays a role in the discrimination of interior states for agents. In some models of agency, input and output data are deliberately distinguished from interior states. Here, we assume that any registers that represent ports for communicating with other agents belong to the interior states of an agent, differing only in their promises. Any value which has arrived from outside, or which is ready for sending, is on the interior of the agent. Sometimes it is expedient to partition interior states by various criteria, creating a ‘split brain’ model [31] based on differing promises, but for now we simply group all the states under the set q_A , for agent A.

Agents are thus defined in part by their boundaries, which are assumed to be finite and closed. That is trivial for elementary agents, but it becomes more important when we compose agents into superagents and scale promises [30].

Definition 1 (Interior and exterior). *The interior of an agent includes its boundary, and the exterior is everything beyond that.*

When we characterize the meaning of ‘intentions’ for agents, we are talking about the possible process outcomes of the agent. Boundaries thus divide intent into precisely two kinds.

Lemma 1 (Interior exterior split). *The process outcomes of any agent must concern either interior (self) or exterior (non-self). There are no other possibilities.*

The expression of intended outcomes for interior (self) are what we call Promises of the First Kind. Intended outcomes hoped or induced by an agent for its exterior (non-self) are called impositions.

In general, the intention to induce outcomes on the exterior of an agent is useless or at best uncertain, since the principle of autonomy or locality implies that agents only control their interior resources with their own processes.

Promises of higher order (second and third kinds) refer to exterior agents, but may be expressed in terms of impositions and promises of a cluster of agents, so first kind promises are the primitive. This is summarized by the first principle for Promises of the First Kind:

Principle 1 (Autonomy). *An agent may not make promises (of the first kind) on behalf of an agent other than itself, that is promises derive from its interior.*

We state without proof here that any interaction between entities can be represented by agents, promises, and impositions [4].

3.2 Absorption and emission of agents

As we scale the composition of agents, it must be possible for agents to be ‘absorbed’ or owned by a superagent by moving from the exterior of the boundary to the interior. An imposition to absorb an agent has the semantics of an attack. A promise to join has the semantics of joining or membership. Imposition is thus a prototype for attack.




Figure 3: Agents can be absorbed and emitted by superagents that compose them. In (a) A and A' are independent and cannot make promises about each other, in (b) A' becomes part of A (self) so that A can now autonomously make promises that collectively include A' .

4 Promises, impositions, and processes

A promise, written:

$$\pi_S(b) : S \xrightarrow{b} R, \quad (2)$$

should not be confused with a morphism or mapping. Rather, it is to be understood as an intended constraint b on the behaviour of the originating agent S (known as the promiser), shared with one or more recipients R , and with body constraint b . At this stage the value of an arrow notation seems spurious at best, but – as we shall see – it unfolds into a powerful description of the directedness of complex causality. The promise body b may be used to express alignment over

its domain. Two promises with bodies b and b' may be aligned according to the overlap of their bodies: $b \cap b'$.

4.1 No determinism implied

It may be difficult for some readers not to associate arrows with the implication of strict determinism – a habit that comes to us from a long tradition of abstracting ‘flow’ in machinery and mathematical constructions (for example vectors, morphisms, etc) that express algebraic and logical certainty. However, we emphasize that this should not be assumed the case about promises. Arrows represent a direction from source to receiver that concerns potential alignments of process outcome, but there is no suggestion of immediacy, completeness, nor absoluteness in this alignment. When we throw someone a ball, it is not guaranteed to be caught, whether the receiver promises to catch it or not. When we throw a series, perhaps only catching half might be promised.

A promise body typically has a type and a constraint on outcome (a goal), representing intent within the same language of sets and subsets. So roughly speaking, a promise for an agent A is something that associates it with a particular outcome expressed in terms of its internal states:

$$“\pi_A \sim A \xrightarrow{\cong} q_A” . \quad (3)$$

The subscript indicates that π_A is a constraint only on A , as otherwise it would be an ‘imposition’, to be described below. Our goal is to be more precise about this association, and – while straightforward – this is where the particular benefits and subtleties of Promise Theory lie.

A promise is a declaration of an intended outcome, like a fixed point. Without processes to bring about the outcome, promises would be toothless empty data. A process, represented as a sequence of steps (that is intermediate interior promises made by agents formed from the interior states) might be called an implementation of the promise as a progression of sub-promises. Promises can exist on a number of levels, as long as we can define the agents that make them.

4.2 Offer, acceptance, and distinguishability

One detail of great importance is that the bodies of promises which make offer or advertise an agent’s service are labelled $+b$ and are often referred to as $(+)$ promises; conversely, those which accept or make use of $(+)$ promises are labelled $-b$ and referred to as $(-)$ promises. Both are needed in order for influence to propagate between agents.

If a collection of agents makes identical promises (called $+b$), then they are indistinguishable to agents that accept everything equally from all $-b$. However, agents may still accept those promises differently, accepting say $b_i \subseteq b$ and thus discern a distinction by selection that was not offered from the source. Interactions are always bindings of this nature, involving $+$ and $-$.




Figure 4: Schematic of relationship between processes on the exterior of an agent and the interior processes that assess, evaluate, and select possible new promises based on receipt of dependencies. The response selection is an autonomous process, so we emphasize that one cannot induce cooperation without remote controlling this process from outside – which violates the principle of autonomy. This is usually taken for granted in ballistic models of reason and action.

If the same agents promise a variety of offered outcomes, we now have a choice of representation: either we say that there are indistinguishable particles, which can be in several states at the same time (for example a superposition, or melting lattice etc), or that the distinction is what identifies them. Both conventions seem to be in use in descriptions of phenomena. For example, when particles are assumed to be embedded in a coordinate system, all with different coordinates, then the coordinates can be said to distinguish them; alternatively, one could say that the position of the particle is indeterminate by symmetry, and position is delocalized. This is the confusion of labelling and observability that riddles all information science.

4.3 The meaning of intent and promise

In order to bridge the differences in modelling viewpoint, and show that they are complementary rather than at odds, we can detail promises and impositions as a series of maps between the different types of structured information – by focusing on what is executed by processes that seek to implement the intended outcomes (keep promises, etc). Such a ‘categorical’ view probably has a limited use, but could help to lay bare the assumptions about representation more clearly for those with an interest in functional representations.

Let us examine the structural components of promises in this way. There is a bewildering array of vocabulary in common use for the elements of data and process, so we make some consistent choices, with our intended audiences in mind. State variables $m \in q_A$ are essentially ‘memory locations’ which can promise to represent different values v , for example spin-up or spin-down, 1 or 0, true or false, real, integer, struct, etc.

The functional values, over a domain of states, are collectively referred to as configurations of the states, denoted $Q(q_A)$, when in an arbitrary representation.

So for an agent S , a configuration is locally a map from some representation of states to values:

$$Q_S : q_S \mapsto v_S, \quad (4)$$

and this generalizes over a collection of agents $\{A\}$:

$$Q_{\{A\}} : q_{\{A\}} \mapsto v_{\{A\}}, \quad (5)$$

The agent's name S or A plays the role of a coordinate subscript, since values may be parameterized by position, so that the configurations of a collection of agents, in the role of 'spacetime' have the role of an evolving function $Q(x, t)$ over all the agent locations x . This form is the one most familiar to programmers in computer science.

It is convenient, for future reference to the usual representations of quantum physics, to encode the same information as a vector with one row per possible state value in q_S , so that the rows are purely numerical and may be associated with statistical ensembles, as in quantum mechanics, for instance.

$$\psi_S : q_S \otimes v_S \mapsto [0, 1], \quad (6)$$

In general, the right hand side could be any non-negative value, representing an average weight; this is common in statistical mechanics and quantum mechanics interpretations of the wavefunction [32, 33]. In this form, the vector ψ_S may be trivially decomposed into an orthonormal set with positive coefficients that represent a configuration weight, and the transition matrices are all unitary and adjoint, forming a simple vector space, whose real valued significance is as given as effective statistical measures.

For agents, the space of possible configurations fills the role of possible 'desired states', 'goals', or 'intended outcomes'. There is a simple association between 'Possible Intentions' ι_S and resulting 'Possible Configurations', captured as the domain of ψ_S for each agent S , written $\Psi_S = \text{dom}(\psi_S)$.

A selection from this domain is a map from the complete domain of configurations to a stable subset (in the graph theoretic sense of [34, 35]) of those configurations $\psi_S(q_S) \in \Psi_S$:

$$\text{Selection} : \Psi_S \mapsto (\psi_S \subseteq \Psi_S) \quad (7)$$

$$\Psi_{\{A\}} \mapsto (\psi_{\{A\}} \subseteq \Psi_{\{A\}}) \quad (8)$$

Selections belong to the power set of Ψ_S . So now we can define the meaning of an actual intention as an association between an agent and the selection of a possible outcome configuration:

$$\iota_S : S \mapsto \psi_S(q_S). \quad (9)$$

Note that intentionality has nothing to do with free will; it is just the ability (by some process) to select between possible alternative outcomes, like a compass direction, or an option one might choose from a menu at MacDonald's.

From this, we view intentions as the bodies of promises, that is the body of a promise is to be understood as an intention which has been turned into a message.

$$\text{Body} : \text{Intention} \mapsto \text{Message} \quad (10)$$

$$b_S : \iota_S \mapsto M_S(\iota_S) \quad (11)$$

To complete the notion of a promise, this message need to be shared within a certain scope (written $\sigma = \{A_i, \dots\}$ in the standard notation of [4]), and consisting of a collection of receiver agents. Normally, we single out a particular ‘promisee’ agent and write it as the receiver R as an intended stakeholder in cooperative behaviour. This terminology of stakeholders, promisers, and promisees, etc, is modelled on human concerns, but also works harmlessly on other scales and in other contexts.

$$\pi_S : S \otimes b_S \mapsto \{R, \sigma\}, \quad (12)$$

The faithfulness of all these mappings, for example between intention and message is now in question. There may be differences between:

- What is intended (type and target constraint).
- What can be represented in the language L that contains M .
- What is thus promised.
- How the promise is received and understood by agents in scope.
- The resolution with which agents in scope promise to build on their understanding of the message.

Which is these can be said to ‘be’ the intent of the agent, and according to whom? These are non-trivial issues, but are usually swept under the rug of convenience.

Sometimes there are special agents to whom promises are directed, because they will benefit from the outcome, but that does not preclude other agents from gaining knowledge of the promise. In physics, just as a specific test charge might be the intended recipient of an electric field’s promise of influence, the same field promise may also be visible to other charges whose future behaviours may also be affected by accepting the field’s influence, whether of interest to instigator of the field or not. Here the scope becomes the sum of the agents, and the intention maps only to the promiser(s).

4.4 Imposition

An imposition is an attempt to induce an outcome in an agent other than self (which – being impossible in general – has far less certainty than promises about self). Impositions can assume without evidence capabilities of a remote agent, because no agent has direct knowledge of another unless that has been promised

(as well as accepted and understood). Promises to accept in advance of impositions can declare a remote agent's capabilities, allowing them to align anyway – thus imposition works best in a framework of promises. Impositions are written:

$$I_S : S \xrightarrow{+b} \blacksquare R. \quad (13)$$

The block arrow is supposed to remind us of a fist. Here the sender S is trying to induce a change in the receiver R . This is not possible without its cooperation, in the shape of a promise to accept the imposition:

$$\pi_R^{(-)} : R \xrightarrow{-b} S. \quad (14)$$

A body constraint in an imposition, by S now refers to a target selection of outcome in R , not S ,

$$\text{Selection} : \Psi_R \mapsto (\psi_R \subseteq \Psi_R) \quad (15)$$

However, this selection is only wishful thinking; S has no control over R . An imposition is only a suggestion. Only the promise to accept the imposition in (14) can result in this selection being made by R , because agents are assumed autonomous. The imposition body is then the transformation of this intent into a message, as before. Note that, within the assumptions of Promise Theory, the only way for an exterior agent to force itself upon another agent would be for it to completely absorb the agent so that it became part of the enveloping agent's state space, for example like taking over a company. Finally, the imposition itself is the association of the imposer S with an outcome in the imposee R , which is more complicated than for a promise:

$$\text{Imposition} : S \mapsto (R \mapsto \text{Selection}) \quad (16)$$

The assessment of success for an imposition is not based on first hand information, as is the case for a promise; it relies on feedback from a remote source – which in turn has to be promised and trusted.

4.5 Signals and messages

Signalling is a well understood term, meaning a kind of message. A message is a (non-local) process that happens on the exterior of the agent boundary for each scale, forming a message channel [36, 37]. In the scaling of this scenario, messages can naturally be passed between the sub-agents of an agent on the interior, etc. The nature of the channel need not be defined. An agent signals its intent, as a promise, by encoding that intention as a message. The nature of the message is not defined a priori. All that matters is that it is encoded in some language $L = \Sigma^*$ with alphabet Σ , that can be received and comprehended by other agents.

As usual, we assume that no signal propagates without acceptance, that is both

$$S \xrightarrow{+b} R \quad (17)$$

$$R \xrightarrow{-b} S \quad (18)$$

are needed for the message in b to be passed from S to R . Further, an agent that cannot comprehend or resolve Σ or L is incapable of accepting a promise.

It is up to each agent to promise what it does with each message that it accepts and interprets. Causality, however, implies that messages are dependent on some function of the interior state of an agent. If messages are functions of an agent's state, then for a promise:

$$\pi_S^{(+)} : S \xrightarrow{+b} R, \quad (19)$$

from sender S to receiver R ,

$$b = \langle \tau, \chi \rangle \quad (20)$$

$$= M(q_1, q_2, \dots) | q_1, q_2, \dots \quad (21)$$

where τ represents the type and its maximal set of possibilities, and χ is a restriction to a subset of promised outcomes. Ultimately, originating within the agent, this has to come from $q_1, q_2, \dots \in q_S$, that is each message is map from the interior states of an agent to a set of strings Σ_S^* forming an interchange language L_S [30, 38, 39]. Note that what can be promised by an agent may be limited, in practice, by the state space of the agents. One cannot represent a real number with a finite number of states, for instance.

Once a promise has been advertised, processes are implied to keep the promise. In the course of promise keeping, a multitude of other messages may be passed that are subordinate to a given promise, and in that sense the fact that they are new promises is of secondary interest and we can simply refer to them as messages, data, or interactions.

Since each agent is causally independent, a priori, interactions between agents occur non-deterministically. Agents can order their own promises and internal processes. The deterministic composition of clocks belonging to a mutually ordered collaboration between agents A and A' is a function of the entanglement of both [31], but remains non-deterministically related to other agents.

Causality is defined in terms of the order of changes in a system.

- Exterior changes are messages.
- Interior changes are processes.

When we rescale a system, the status of processes and messages can therefore be interchanged.

4.6 Messages between agents

Messages are the only way for agents to share information and induce action. A message is a string of symbolic information, in the Shannon sense, that is emitted by one agent S and which may be received and accepted by another agent R [36,37,40]. Messages imply no functional outcome unless they are accepted by a receiver. Messages that declare a promise are (in principle, if not in practice) distinguishable from messages involved in the keeping of promises. This may only be a question of representation.

Messages $M(q_A)$ can only be based on the configuration of state(s) of an agent $\psi(q_A)$ at each moment:

$$M : \psi(q_A) \mapsto \Sigma_A^*. \quad (22)$$

and are therefore a map from those states to strings of an alphabet used in communication. Messages may or may not be promises, but all messages are composed serially (sequences) or in parallel (vectors).

Computer science and physics alike often want to think in terms of actions (at least at a classical level) when, in an information picture, the only actions are messages between interior processes. All physical quantities (like quantum numbers and momenta) are passed by messages, or the composition of messages. One may distinguish different types of message, for convenience, but there is no need to any other kind of influence. The reorganization of agent ownership that occurs when we pay for something with a dollar bill and take away our goods is simply an illusion of scale. On a higher level, such abstractions might be useful as shorthands for compositions of promises and impositions, but they have no place at an elementary level.

5 Keeping promises

Agents can only make Promises of The First Kind about their own states and selections, and – when promised observable by other agents – these can be assessed by those remote agents. The remote agents, in turn, can only keep their own promises about their own resources, which includes the promise to accept and assess the promises offered to them using their own interior resources. This restriction makes obvious sense when applying to phenomena on a microscopic level (the absorption of a photon, say, would depend on both offer and acceptance, or ψ^+ and ψ^-), but it turns out to be equally revealing when trying to establish dynamical models of socio-economic phenomena too [21].

Promises can therefore only be kept by processes within the agent that makes the promise. In a binding:

$$S \xrightarrow{+b} R \quad (23)$$

$$R \xrightarrow{-b} S, \quad (24)$$

the process that keep the promises are on the left hand side of the arrows. The right hand agents are passive observers. Note, however, that the acceptance (-)

of the offer promise (+) is downstream of it, and therefore the propagation of intent ultimately depends on the keeping of the (-) promise in a binding.

A transition $\Delta\psi(q_A)$ is a ‘change of configuration’ meaning a change in the values of states q_A :

$$\Delta\psi : \psi \mapsto \psi' \quad (25)$$

A transition that keeps a promise is one in which the end state is an intended state:

$$\Delta\psi : \psi \mapsto \psi' = \iota(S). \quad (26)$$

For an intention to propagate, this requires promise keeping from both sides of a promise channel (source and receiver), comparable to the mutual information [40]:

$$\psi_S(q_S) \mapsto \iota_{\pi^+}(S) \quad (27)$$

$$\psi_R(q_R) \mapsto \iota_{\pi^-}(R), \quad (28)$$

where ι_{π} are the intentions of the agents to bind matching promises π^+ and π^- . The strength of the binding is $\psi_S \cap \psi_R = \iota(S) \cap \iota(R)$ (see Figure 5). The




Figure 5: A promise channel is an overlap in the intentions between a sender S and a receiver R . The overlap is communicated over a message channel, in the normal sense of a Shannon communication channel.

extent to which a promise is kept on a statistical basis can be assessed using the mutual information of the assessments of the channels [40].

Processes can also be represented as promise graphs, using conditional promises to define the causal flow of information, for example a transition within a superagent can also be represented by a chain of promises between its subagents:

$$A_{in} \begin{array}{c} \xrightarrow{+b_0} \\ \xleftarrow{-b_0} \end{array} A_1 \xrightarrow{b_1|b_0} \dots A_{out} \quad (29)$$

6 Assessment

In order to assess a remote agent S 's promise, an agent R can only use the information it has been promised (and which it has accepted and received). So it must accept messages arriving from S , and promise to represent these in its own interior states q_R .

In the general case, an agent makes and alters its promises based on:

- *Internal information held in its own internal state q_R ,*
- *Its assessments of promises $\alpha_R(\pi_S)$ of counterparts S .*

An assessment is one of many processes that an agent can use as a basis for decision-making (selection). Assessments are essentially arbitrary and subjective mappings from a given state to some outcome. Assessments can only be made based on information within an agent (which includes information promised by other agents and sampled before the assessment).

- *A rational promiser will assess a promise to be kept if the end state of the promise-keeping trajectory coincides with the declaration of the outcome.*
- *A rational promisee will assess a promise to be kept if the consequences for its own interior states reach its interpretation of the desired end state.*

Note that these statements also apply to (-) promises to accept another agent's promise. Without processes, the existence of the promise could not be received, sampled, detected, parsed, or comprehended. Without processes, there can be no messages and no transitions.

6.1 Derivative obligations

In classical logic, obligations are assumed to be the primitives and promises are assumed to induce obligations to act. This view is a hangover from a tradition of thinking in terms of ballistic causality. In Promise Theory the structure is opposite, in keeping with modern information-based ideas of force transmission, especially from the quantum theory. We can define an obligation as follows.

Definition 2 (Obligation). *A positive assessment α_R by an agent R about the need to make a promise $\pi_{R|S}$, based on its acceptance of a prior imposition ι_S or promise π_S , from from an exterior agent S .*

The obligation is an assessed impetus (also called a valuation in [4]) for a new conditional promise which is conditional on π_S or ι_S . We therefore see why it is common to associate obligations with a (relativistic) sense of values.

Suppose an agent S accepts a promise from a third party T . An obligation for S to promise $+b$ to some agent R (possibly equal to T) is the assessment $\alpha_S()$

(made by S) that a promise is has received should cause it to make a promise of its own conditionally on assessing that the promised dependency D is received:

$$\pi_S : S \xrightarrow{+b \mid D} R \quad (30)$$

$$S \xrightarrow{-D} T \quad (31)$$

The obligation is thus self-induced by the autonomous self-assessment of the agent S . It is not induced from outside – thus maintaining locality.

The attempt to impose an obligation, by sending a proposal of intent to be implemented by a receiver, would look like this:

$$S \xrightarrow{+\text{def}(t)} R. \quad (32)$$

Note, however, the error in this thinking is the belief that the imposition must be accepted and implemented by R , at the best of S . As noted above, only R can decide about R , and only conquering the agent could enable S to achieve that. The illusory belief in the effect of obligation is a view that presupposes alignment with a set of standard behaviours, which is unwarranted (though perhaps more common in the era of its origin).

6.2 Example: Non-deterministic State Machine

Let us briefly compare a simple graph, as one might find as part of a non-deterministic automaton (see Figure 6). In the classical picture, which is based on ‘machine thinking’, a programmer or designer can specify transition rates (probabilities) for transitions between the nodes. In Figure 6(a), it is assumed that as soon as a signal or token arrives at the first node A_1 , it ‘must immediately’ trigger a reaction that continues in the direction of the arrows, first to A_2 , and from there one of two things can happen: either an immediate transition occurs to A_3 or to A_4 with probabilities 0.7 and 0.3 respectively. These probabilities must sum to 1, meaning that there is no possibility that a transition would not take place.

In Figure 6(b), we start from a configuration in which each agent makes promises already, similar to the already-designed state machine, but now we see the doubling of arrows with polarities $+$ and $-$ for what is offered and accepted between the agents. These promises imply no timescale; indeed, the response of each agent is entirely autonomous, and at its own behest. Each promiser A_i promises to pass on something $+m_i$ and the promisees promise to receive something $-m_i$. The overlap $+m_i \cap -m_i$ is not implied: each agent makes this decision independently. We are no longer able to specify the distribution of probabilities for passing on a message, or whether the choices will be mutually exclusive (without additional promise constraints), since we cannot say whether what is offered will be accepted. This is not entirely up to the local node anymore – it is actually up to the receivers in the final instance. Each promisee (of either

+ or -) can assess whether it considers the promises were kept:

$$\begin{aligned} \alpha_2(+m_1), \alpha_1(-m_1), \alpha_2(-m_2), \alpha_2(-m_3), \\ \alpha_3(+m_2), \alpha_4(+m_3), \text{ etc.} \end{aligned} \quad (33)$$

In order for the situation in (a) to be reproduced, we might expect to have that

$$\alpha_2(-m_2)\alpha_3(+m_2) \simeq 0.7, \quad (34)$$

however, this now a non-local assessment. It is not available to any agent in the graph. What we normally expect is that there is a godlike observer agent G with access to all of the promises in the state machine, and who can assess these probabilities itself, according to a common standard (calibration).

$$\alpha_G(-m_2)\alpha_G(+m_2) \simeq 0.7, \quad (35)$$

In other words, we cannot say that the response is immediate or even that the probabilities will sum to 1 at each node. This is the picture assumed in quantum mechanics of the Schrödinger equation for instance. However, such a situation is not easy to engineer – as we now know from distributed computing.




Figure 6: A comparison between a conventional ballistic (‘push’) automaton and a promise theoretic model. The distinction concerns where the probabilistic weights for transitions can be asserted. In PT each agent controls only its own behaviour, so even probabilistic transition rates cannot be determined for the entire system – they are only emergent. Classical FSM is unclear on how probability is defined.

7 Propagation and ‘wiring’ in spacetime

How signals and outcomes move through space (and at what rate, for each observer) is the question that enables the building of systems from components on a variety of scales. This has a renewed importance in cloud computing, for instance. Although the issues are quite universal, we confront them in different ways at different scales. Propagation involves not only promising and accepting

but the cooperation of individual agents, and a certain homogeneity in their local assessments. Only then can we form channels for information to pass along reliably and maintain an illusion of homogeneous ‘order’.




Figure 7: Promise trajectories take two main forms i) conditional ‘path integrals’, and ii) convergent fixed points (stable subgraphs), well known from configuration management technology [2, 41].

7.1 Transition functions

There are two main kinds of state transitions (shown in Figure 7). The usual kind of trajectory in Figure 7 a) is the classic ballistic path followed by a linear conditional process, or causal set path. This is a path through state space q , but with conditional promises over a scaled set (in which agent promises become the states in a superagent), this can also be a path through a set of agents (that is a spacetime path).

$$\Delta\psi_{12}|q_1\rangle = |q_2\rangle \quad (36)$$

A second class of states is the so-called stable subgraph or convergent states that are fixed points q_p of a class of convergent transitions:

$$\Delta\psi_{np}|q_n\rangle = |q_p\rangle \quad (37)$$

$$\Delta\psi_{np}|q_p\rangle = |q_p\rangle. \quad (38)$$

Any process trajectory is an ordered composition of transitions operators, mapping from an initial state configuration to a final state. If we use the useful Dirac bra-ket shorthand notation for states $\psi^+ \mapsto |\psi\rangle$, then a spacelike trajectory has the form:

$$|\psi_{out}\rangle = \Delta\psi(q_n) \circ \dots \circ \Delta\psi(q_1) |\psi_{in}\rangle \quad (39)$$

that is the path ordered composition of changes from agent location to agent location. A timelike trajectory, takes the form:

$$|\psi_{out}\rangle = \Delta\psi_n(q) \circ \dots \circ \Delta\psi_1(q) |\psi_{in}\rangle \quad (40)$$

that is a sequence of changes on the same set of states. From the perspective of a physics of promises, there is a lot more to say, but it strays beyond the natural scope of this paper, so we defer this for another time. The patterns in (36) and (37) represent the main cases for retarded and advanced propagation respectively [5]. The body of a promise might consist of a detailed path or merely a desired end state, and what the receiver accepts might be only a subset of this offer. This makes a promise theoretic interaction different from what is conventionally assumed in physics, where conservation laws insist on the equality of offer and acceptance during transactions. The ‘accounting’ status of what measures are promised by an agent is not defined a priori. Whether one should assume conservation of the promised ‘token’ exchange is an open question for each particular formulation.

It is surely worth a brief mention that, if one assumes a description based on quasi-infinitesimal changes (say, by arguing for sufficiently large-scale statistical coarse-graining as we do in physics), then there will be conservation of accounting measure within each agent, allowing transitions to be formulated in the usual path integral representation of a partition-transition function at each location along time-like trajectories with interior conservation along the path:

$$\Delta\psi(q) \simeq \xi_\pi T_{\delta\pi}, \quad (41)$$

where the matrix $T_{\delta\pi}$ is the generator of a transition for a step $\delta\pi$ towards keeping the promise π . There is no basis for making this assumption along spacelike trajectories. So, by the usual technique of exponentiation, $1 + \xi_\pi T_{\delta\pi} \sim \exp(\xi_\pi T_{\delta\pi})$ for infinitesimal:

$$\Delta\psi|\psi\rangle = e^{\xi_\pi T_{\delta\pi}}|\psi\rangle, \quad (42)$$

analogous to the generators of a canonical group [32], and adding a boundary condition constraint on the allowed paths from the current state the path count, which approximates a density of collective path states is

$$\ln \int d\xi_\pi e^{\xi_\pi T_{\delta\pi} + G_\pi}, \quad (43)$$

for the unspecified generator G_π of the initial conditions and subsequent accounting constraints. So transition function ‘amplitudes’ can be obtained in the analogous way as for statistical mechanics or quantum theory, by an effective action of the Boltzmann-Shannon entropy partition form along the paths:

$$\delta\langle\psi_{out}|\psi_{in}\rangle \sim \delta \ln \int d\xi_\pi e^{\xi_\pi T_{\delta\pi} + G_\pi}. \quad (44)$$

This amounts to applying ordered sequences of transition matrix operators $T_{\delta\pi}$ to the composition of micro-transitions that keep promises statistically.

In microscopic physics, one assumes a multitude of conservation principles that account for energy, momentum, charge, etc, which is equivalent to assuming that (+) promises are always accepted (-) fully with some probability. This basically assigns joint non-local responsibility for promise keeping in physical systems in order to maintain conservation or unitarity. Promise Theory does not require this; indeed, on a larger scale it is unlikely to be true unless new constraints are added at each new scale. So the method of trying to count and predict outcomes as probabilities is unlikely to be a simple matter unless one can constrain a system with basic conservation principles. Thus, while fine for the cases considered in microscopic physical systems, Promise Theory proposes to extend interactions to larger scales where no such constraints are in play – such as one might hope to formulate for the socioeconomic sciences. Now the question of interest is not only what is the nature of the beginning and end states, but also whether one can argue infinitesimal transformations of an invariant path function. In the case of monetary economics, it is indeed normal to assume the conversation of monetary tokens, but promises of a broader nature have no such requirement – and thus the dream of a purely calculable algebraic approach for social disciplines (in an Asimov sense [42, 43]) seems not to be realizable without the unwarranted assumption of average conservation of intangibles.

7.2 Clocks and interior time

While agents define space, processes define distributed clocks and therefore the meaning of agent time. An agent’s own perception of time, by sampling, is what we may refer to as its proper time, mapping to the corresponding concept in Einsteinian relativity.

Messages cannot be received without active processes that actively sample shared states at a Nyquist frequency. So the observation of exterior time is the non-local expression of an interior time, which remains unexplained in Promise Theory. This is a ‘cognitive’ picture of observation, which matches with the modern understanding of brain function, as well as basic Information Theory. This is not a vacuous observation, since any causal determinism of interactions between agents is now extinguished by the unknowns of local clock rates.

Promised interconnections of agents compound the complexities, because agents can now use their promised access to other agents to store memory via exterior promises ‘stigmergically’. This means that processes are not limited to the interior of agents, but spill out into a larger scale, forming superagents. Even agents, not formally identified, such as pro-forma agents that represent an ‘environment’ can be used as memory – for example in stigmergic cooperation between agents, employing shared resources, broadcasting, etc. This points to the need for substitute principles to understand what the equivalent of ‘equations of motion’ might be for socioeconomic systems on an agent level. Although a lot of speculative work has been proposed in connection to cellular automata and emergent phenomena, there is no candidate as far as we know today.

8 Summary

Promise Theory concerns the alignment of agent capability and behaviour, in the sense of describing a measure of intent which combines dynamical (quantitative) and semantic (qualitative) factors, and allows one to consider the propagation of influence through a network of possibly hierarchical interactions. Its focus is on the combined roles of these semantics and dynamics in the interactions between causally independent, that is autonomous agents (see table 3). As such, it doesn't equate to any other theory of interaction in detail, although it resembles different aspects of several.

It is natural to ask what the approximate correspondence is between the Promise Theory picture of dynamics and semantics and other theoretical frameworks. Most high level descriptions of process focus on qualitative, symbolic, or semantic (functional) relationships and behaviours. Physics, on the other hand, singles itself out as the one science in which one tends to suppress semantics in favour of quantitative measures – this is possible essentially because there are fewer semantic distinctions of importance at a low and universal level of description. That, in turn, makes the ledger of counting things a simpler matter. Some straightforward though approximate correspondences are illustrated in tables 1 and 2.

PHYSICS	PROMISE
FIELD/CHARGE	PROMISE/TYPE
INTERACTION	PROMISE BINDING
PARTICLE	EXTERIOR PROCESS
VIRTUAL PARTICLE	INTERIOR PROCESS
ENTANGLEMENT	CO-DEPENDENT CONDITIONALS
STATES $q(x, t)$	STATES q_A
CONFIGURATIONS $\psi(x, t)$	CONFIGURATIONS $\psi(A)$
SPACE x	AGENTS A
TIME t	EXTERIOR TIME
FLUCTUATION	INTERIOR TIME

Table 1: The heuristic correspondence between dynamics and semantics in cell biology and in a Promise Theory. It is an open question whether this correspondence can be made precise by a natural transformation in Category Theory.

In spite of the adherence to locality or autonomy, we cannot avoid non-local concepts whenever there is interaction. At a high level, the concept of trust is the natural expression of the quantitative affinity in binding. It is an inevitable side-effect of assessment, which affects the degree or likelihood that

PHYSICS	PROMISE
PROTEINS	PROMISE/TYPE
RECEPTOR BONDS	PROMISE BINDING
CELLULAR PROCESS	EXTERIOR PROCESS
MOLECULAR PROCESS	INTERIOR PROCESS
INTERIOR SYMBIOSIS	CO-DEPENDENT CONDITIONALS
EXPRESSIONS $q(x, t)$	STATES q_A
CONFIGURATIONS	CONFIGURATIONS $\psi(A)$
CELLS / MOLECULES x	AGENTS A
CELL REACTION t	EXTERIOR TIME
CELL MORPHOLOGY	INTERIOR TIME

Table 2: The correspondence between dynamics and semantics in physics and in a Promise Theory. It is an open question whether this correspondence can be made precise by a natural transformation in Category Theory.

autonomous agents will align their intentions and form a promise-binding. Unlike the conventional representation of charge in physics, trust must be a purely local assessment however, so there strict observer relativity is not violated.

There are many issues to pursue concerning the foundations of Promise Theory, but for the present we note that it remains an eminently applicable model of interaction that survives the test of scaling in a way that cannot be said of many models of the natural or technological world.

9 Appendix: Morphisms, Categories, Promises

Promise Theory's successes have been mainly in the engineering sphere, where few mathematically disciplined papers are written – its principles are easily grasped without extensive formalism, and may be easily transcribed into computer code or language directives [5–19]. Category Theory, by contrast, has a large and growing following, despite its expanding jargon and runaround definitions – yet there is ample evidence that Category Theory will increasingly find its way into engineering, where mathematicians roam. (As a side remark, there seems to be no easy way into Category Theory from the outside, as its construction is both recursive, and highly technical. We have not found any one introduction to suffice, but can recommend the articles at [44, 45]. Category Theory's legitimacy is no longer in question as a subject, but one has the sense that practitioners continue to struggle to find pedestrian uses for its ideas to realize the belief that Category Theory will be the framework that subsumes all others.) Lately, we have observed it being used to describe the kinds of static and

	PROMISE	IMPOSITION
NOTATION	$S \xrightarrow{\pm b} R$	$S \xrightarrow{\pm b} \blacksquare R$
ACCEPTANCE	$R \xrightarrow{\mp b} S$	$R \xrightarrow{\mp b} \blacksquare S$
DETERMINISTIC	No	No
REFERS TO	S (self)	R (non-self)
PREREQUISITE: FOR OUTCOME	NONE	$R \xrightarrow{\mp b'} S$ where $b \cap b' = \emptyset$
ASSESSMENT	$\alpha_S()$	$\alpha_R()$
TRUST DEPENDENCE	in R	in S AND R

Table 3: Contrasting promises and impositions.

dynamical systems for which Promise Theory was developed. Given the contortions that are needed to express even simple ideas, it is unclear to what extent the mathematical precision of categories can help at an engineering level. The topic of monads, in particular, is often drawn on as a justification in connection with computer programming.)

Unlike Promise Theory, Category Theory looks for precise analogies, and tends to fall back on existing forms of description, such as topology, geometry, differential equations [26], and state machines [46], subordinating its analogies to those dynamical descriptions. Promise Theory, on the other hand, represents systems on its own terms – emphasizing its few axioms (for better or for worse), and in so doing has some passing resemblances to quantum theory, for reasons we cannot go into here. In so doing, Promise Theory has more of the feel of physics than mathematics. Promise Theory has far fewer concepts, but is also under-specified, meaning there is plenty of room for fusing it with other methods of description, including categories. Its rejection of logical modalities such as obligations is perhaps most noteworthy. Thus, while Promise Theory and Category Theory are very different animals, there is clearly overlap when both frameworks purport to be of such a general nature, and it seems worth a few notes on the bridge between them.

Crudely speaking, Category Theory is a theory of ‘types’, their composition, and the classification of objects into those types, by mappings called morphisms (abstracted from the classical homomorphism, isomorphism, etc). Unlike promises, such mappings are not usually considered to be voluntary acts of cooperation, because the objects in a map are not considered autonomous agents. Morphisms are thus ‘presumptuous’ in the view of Promise Theory – relationships imposed from without rather than aligned by voluntary agency. The types are not so much agents as bricks in an interior ‘game’ of association. Nevertheless, we can represent morphisms as voluntary cooperative structures. The

suggestion that maps can be interpreted the other way around (as cooperative outcomes, is a commonly presumed (though unjustified) idea in Category Theory. By representing this in Promise Theory, we can see the necessary and sufficient basis for making such an identification.

A promise binding is not to be understood as a mapping whose co-domain is guaranteed by an arrow from the domain. If we take a statement like

$$S \mapsto R, \quad (45)$$

and interpret it in terms of promises, the intent of a map is to imply that elements of S ‘promise’ something unequivocally to elements of R , and that objects in R accept these, in turn, unconditionally. In Promise Theory, this assumption would be a violation of the assumption of autonomy of R . In Promise Theory, each agent is in a position to make promises only about itself, respecting what one would call ‘locality’ in physics, but no agent in general has the ‘authority’ to determine global statements that relate to others than itself. At best, a third party observer could promise its own assessment of a relationship between S and R . This is an important detail not to be swept away in our enthusiasm for type relations. In physical terms, it is a recovery of ‘locality’, such as one sees in the shift from Newtonian to Einsteinian relativity.

Suppose we represent a domain of a map as the ‘sender’ agent of a morphism S , and the co-domain as a recipient agent R . The sender superagent promises an association called $+f$ (say a function of S) to the co-domain R , which acts as recipient. This only works in Promise Theory if the image superagent R accepts the offer completely, by promising exactly $-f$ in return. We can now claim that the promise has been offered and accepted between the domains in their role as agents, or vice versa.

The relation is on the scale of domain agents, but we might also need to know the kind of map by relating the elements, which are the interior states of S and R . An object of the domain q_S is a state within agent S , and it maps to an image q_R in R if and only if there is a state q_R which promises to accept the value v_S promised by q_S , and that value is accepted precisely by the image state q_R . Having received these data, the agent R is now in a position to make a conditional promise that its own interior states represent the conditional result $f(S)$ given S , which is written $f(S)|S$, for each state member.

$$\Pi = \left\{ \begin{array}{l} S \xrightarrow{+f} R \\ S \xleftarrow{-f} R \\ q_S \xrightarrow{v_S} q_R, A? \\ q_R \xrightarrow{-v_R} q_S, A? \\ R_j \xrightarrow{q_R=f(q_S)|q_S} A?, \end{array} \right. \quad (46)$$

A mapping f between autonomous domain objects S and autonomous co-domain objects R would thus correspond to a third party assessment made by an unspecified observer $A_?$:

$$\alpha_{A_?}(\Pi) : 'R = f(S)' \mapsto \text{True}. \quad (47)$$

Normally, however, the mapping is in an abstract space of possibility, realized by some more physical object like a computer ψ , which would then promise a function declaration

$$\psi \xrightarrow{+f:S \mapsto R} A_? \quad (48)$$

or

$$\psi \xrightarrow{R=f(S)} A_? \quad (49)$$

The importance of this construction in *Promise Theory* is that it adds a real-world relativity and causal uncertainty into the presumed association between objects, which seems to be entirely un-modelled with categories. It doesn't automatically require properties of categories or groups, for instance, without explicit declaration. Ordinarily, logics and categorical relations have the status of impositions on representative objects [4]. Concepts like operator and function are all built on imposition.

The use of wiring diagrams between 'machine representations' in *Category Theory* feels superficially similar to the concepts used in *Promise Diagrams*, modulo the binding of $+$ and $-$ promises implied by locality (autonomy), but there are significant differences that we won't go into here. There are equivalents; for example, variable-substitution wires between 'resource-sharing machines' [47] that have been modelled as 'decorated co-spans' [48] in *Category Theory* behave like the 'matroid constructions' in *Promise Theory*, and the bi-directional 'lens' construction [49] suggest a similarity to *Promise Theory's* bidirectional promise patterns, but they represent a quite different circularity without the requirement of commuting morphisms. There are many interesting avenues one could explore between the promises and categories, but we defer that for another occasion. Exposing these matters in detail would be an interesting exercise for someone with time on their hands!

References

- [1] Mark Burgess. *An Approach to Understanding Policy Based on Autonomy and Voluntary Cooperation*. In IFIP/IEEE 16th International Workshop on Distributed Systems Operations and Management (DSOM), in LNCS 3775, pages 97–108, 2005.
- [2] M. Burgess. *A Site Configuration Engine*. Computing Systems (MIT Press: Cambridge MA), 8:309, 1995.

- [3] M. Burgess and R. Ralston. *Distributed Resource Administration Using CFEngine*. Software Practice and Experience, 27:1083, 1997.
- [4] J.A. Bergstra and M. Burgess. *Promise Theory: Principles and Applications* (second edition). *χtAxis Press*, 2014,2019.
- [5] M. Burgess. *A Treatise on Systems: Volume 2: Intentional Systems with Faults, Errors, and Flaws*. ISBN-13 : 979-8603300092 (2020). <http://markburgess.org/treatise.html>
- [6] S. Hogg. *Promise Theory—Can You Really Trust the Network to Keep Promises?* Network World, 2015. <https://www.networkworld.com/article/2449562/promise-theory-mark-burgess-cfengine-sdn-cisco-aci-apic-opflex.html>
- [7] Unattributed article. *Why you need to know about promise theory*. No Jitter, 2018. <https://www.nojitter.com/why-you-need-know-about-promise-theory>
- [8] Zeus Kerravala. *Cisco expands intent-based networking portfolio*. No Jitter, 2018. <https://www.nojitter.com/enterprise-networking/cisco-expands-intent-based-networking-portfolio>
- [9] C. Hemm Klok. *The Quest To Make Code Work Like Biology Just Took a Big Step*. Wired, 2016.
- [10] ‘Melissa’. *Opflex-ing your Cisco Application Centric Infrastructure*. Vmiss.net, 2014. <https://vmiss.net/opflex-ing-your-cisco-application-centric-infrastructure/>
- [11] M. Chandramouli and K. Sivakumar. *Concepts of network intent*. IRTF draft draft-moulchan-nmrg-network-intent-concepts-00, 2014. <https://tools.ietf.org/html/draft-moulchan-nmrg-network-intent-concepts-00>
- [12] A.W. Morales. *The trouble with Promises: Patrick Debois Explains Serverless and ‘Service-full’ Culture*. Forbes, 2017.
- [13] D.C. Marinescu. *Cloud Computing: Theory and Practice*. Morgan Kaufmann, 2013.
- [14] J. Sussna. *Turning Requirements Into Promises*. Blog post, 2013. <http://goodrequirements.com/pt/>
- [15] J. Sussna. *Scaling Agile Projects Using Promises*. Blog post, 2013.
- [16] J. Sussna. *Promise Theory, Devops, and Design*. Blog post, 2013.
- [17] J. Sussna. *Promise Theory, DevOps, and Design*. Blog post, 2013.

- [18] *J. Sussna. Designing Delivery: Rethinking IT in the Digital Service Economy. O'Reilly, 2015.*
- [19] *C. Ciceri. Notes on Domain-Driven Design Europe ii - 2018 Edition. DZone, 2018.*
- [20] *M. Burgess. Thinking In Promises: Designing Systems for Cooperation. O'Reilly Media, 2015.*
- [21] *J. Bergstra and M. Burgess. Money, Ownership, and Agency. χt -axis Press, 2019.*
- [22] *J. Bergstra and M. Burgess. Promise Theory Case Study on the 2016 Brexit Vote. χt -axis Press, 2017.*
- [23] *J. Bergstra. Promises and Threats by Asymmetric Nuclear Weapon States. χt -axis Press, 2019.*
- [24] *S. Eilenberg and S. MacLane. General Theory of Natural Equivalences. Transactions of the American Mathematical Society, 58:231–294, 1945.*
- [25] *J. Bergstra, A. Ponse, and S.M. Smolka, editors. The Handbook of Process Algebra. Elsevier, 2001.*
- [26] *J.C. Baez and B.S. Pollard. A Compositional Framework for Reaction Networks. arXiv:1704.02051 [math-ph], 2017.*
- [27] *G.R. Grimmett and D.R. Stirzaker. Probability and Random Processes (3rd edition). Oxford Scientific Publications, Oxford, 2001.*
- [28] *C. Hewitt, P. Bishop, and R. Steiger. A Universal Modular Actor Formalism For Artificial Intelligence. In Proc. of the 3rd IJCAI, pages 235–245, Stanford, MA, 1973.*
- [29] *J.A. Bergstra and M. Burgess. Local and Global Trust Based on the Concept of Promises. Technical report, arXiv.org/abs/0912.4637 [cs.MA], 2006. <https://arxiv.org/abs/0912.4637>*
- [30] *M. Burgess. Spacetime with Semantics (ii). arXiv.org:1505.01716, 2015.*
- [31] *P. Borrill, M. Burgess, A. Karp, and A. Kasuya. Spacetime-Entangled Networks (i) Relativity and Observability of Stepwise Consensus. arXiv:1807.08549 [cs.DC], 2018. <https://arxiv.org/abs/1807.08549>*
- [32] *J. Schwinger. Quantum Kinematics and Dynamics. Addison Wesley, 1970.*
- [33] *L.E. Ballentine. Quantum Mechanics: A Modern Development. World Scientific, 1998.*
- [34] *C. Berge. The Theory of Graphs. (Dover, New York), 2001.*

- [35] *D.B. West. Introduction to Graph Theory (2nd Edition). (Prentice Hall, Upper Saddle River), 2001.*
- [36] *C.E. Shannon and W. Weaver. The Mathematical Theory of Communication. University of Illinois Press, Urbana, 1949.*
- [37] *T.M. Cover and J.A. Thomas. Elements of Information Theory. (J. Wiley & Sons., New York), 1991.*
- [38] *M. Burgess and S. Fagernes. Laws of Systemic Organization and Collective Behaviour in Ensembles. In Proceedings of MACE 2007, volume 6 of Multicon Lecture Notes. Multicon Verlag, 2007.*
- [39] *M. Burgess and S. Fagernes. Laws of Human-Computer Behaviour and Collective Organization. submitted to the IEEE Journal of Network and Service Management, 2008.*
- [40] *M. Burgess. Information and Causality in Promise Theory. arXiv:2004.12661 [cs.MA], 2020.*
- [41] *M. Burgess. Configurable Immunity Model of Evolving Configuration Management. Science of Computer Programming, 51:197, 2004.*
- [42] *Isaac Asimov. Foundation. Doubleday, 1951.*
- [43] *S. Galam. Sociophysics. Springer, 2012.*
- [44] *T.D. Bradley. Website math3ma blog article series. <https://www.math3ma.com/>.*
- [45] *The n-Category café. Website blog article series. <https://golem.ph.utexas.edu/category/>.*
- [46] *B. Fong and D.I. Spivak. Seven Sketches in Compositionality: An Invitation to Applied Category Theory. arXiv:1803.05316 [math.CT], 2018.*
- [47] *Sophie Libkind. Unifying Open Dynamical Systems: An Algebra of Resource Sharing Machines. YouTube video, 2020. <https://www.youtube.com/watch?v=cjti9KdXLY4>*
- [48] *Brendan Fong. Decorated Cospans. arXiv:1502.00872 [math.CT], 2015.*
- [49] *D.I. Spivak. Generalized Lens Categories via Functors $C^{op} \rightarrow Cat$. arXiv:1908.02202v3 [math.CT], 2020. <https://arxiv.org/abs/1908.02202>*