

The Logic of Typicality

Harry Crane and Isaac Wilhelm

Penultimate draft of Feb 6, 2019.

Please cite the version in Valia Allori (ed.), *Statistical Mechanics and Scientific Explanation: Determinism, Indeterminism and Laws of Nature*.

Abstract

The notion of typicality appears in scientific theories, philosophical arguments, mathematical inquiry, and everyday reasoning. Typicality is invoked in statistical mechanics to explain the behavior of gases. It is also invoked in quantum mechanics to explain the appearance of quantum probabilities. Typicality plays an implicit role in non-rigorous mathematical inquiry, as when a mathematician forms a conjecture based on personal experience of what seems typical in a given situation. Less formally, the language of typicality is a staple of the common parlance: we often claim that certain things are, or are not, typical. But despite the prominence of typicality in science, philosophy, mathematics, and everyday discourse, no formal logics for typicality have been proposed. In this paper, we propose two formal systems for reasoning about typicality. One system is based on propositional logic: it can be understood as formalizing objective facts about what is and is not typical. The other system is based on the logic of intuitionistic type theory: it can be understood as formalizing subjective judgments about typicality.

1 Introduction

Typically, gases in non-equilibrium macrostates evolve to the equilibrium macrostate relatively quickly. Not all gases do: in fact, the initial microstates of some gases prevent them

from ever reaching equilibrium. But those initial microstates are unusual, or atypical. Nearly all initial microstates are not like that. Nearly all initial microstates lead to equilibrium after a short while. This is a ‘typicality fact’: a fact about what is typical.

Typicality facts are studied in many areas of science, but they are particularly prominent in statistical mechanics and quantum mechanics. For instance, Boltzmann discusses a version of the typicality fact just mentioned: the overwhelming majority of initial conditions of a gas, he writes, reach equilibrium in a relatively short amount of time (1896/2003, p. 394). On the basis of his many-worlds interpretation of quantum mechanics, Everett argues that typically, the probabilistic predictions of the Born rule are valid (1956/2012, p. 123). In their analysis of Bohmian mechanics, Dürr, Goldstein, and Zanghì show that typically, initial configurations of the universe lead to empirical distributions that agree with the probabilistic predictions of the quantum formalism (1992, p. 846). Reitmann (2007) shows that given certain generic conditions, pure quantum states typically yield more-or-less the same expectation values for sets of observables which are not too large. Kiessling (2011) shows that for N gravitationally-interacting bodies confined to the surface of a sphere, Boltzmann’s H functional is minimized by states which are typical for those N -body systems (in the limit as N approaches infinity). Tasaki (2016) shows that pure quantum states in the microcanonical energy shell typically share a particular collection of properties associated with thermal equilibrium.¹

These typicality facts are explanatory. The typicality fact about the initial conditions of gases, for example, explains their thermodynamic behavior (Goldstein, 2001, p. 52). Both the typicality fact discussed by Everett and the typicality fact discussed by Dürr, Goldstein, and Zanghì provide explanations of why the frequencies observed in quantum experiments conform to the probabilities predicted by the Born rule. There is, as Goldstein puts it, a “logic of appeal to typicality” in scientific explanation (2012, p. 70).

The logic of typicality extends beyond its role in explanation, however. Typicality is invoked in many different kinds of scientific reasoning: explanation, prediction, evaluation of

¹For examples of typicality results in mathematics, see (Kesten, 1980), (Alon et al., 1998), and (Ledoux, 2001).

hypotheses, and more. But what is the logic of that reasoning?²

Typicality is also invoked in non-rigorous mathematical reasoning. Take Goldbach's conjecture,³ which has been shown to hold for all natural numbers less than 10^{18} .⁴ The set of integers on which the conjecture has been verified is large but finite, and therefore does not pass the 'nearly all' threshold to be considered typical in the sense described above. But consider a number theorist whose past experience is such that when a mathematical claim has been verified on a similarly large sample of natural numbers, that claim has most often turned out to be true. Based on this experience, the number theorist may feel justified to conjecture that the claim holds in this specific case. Such judgments are based on assessments of typicality: they are based on reasoning to the effect that past conjectures which have been empirically tested to the same degree as Goldbach's conjecture have typically turned out to be true. This is a 'typicality judgment': a judgment about what is typical.

Typicality judgments are common in science, mathematical conjecture, and everyday reasoning. They help scientists arrive at hypotheses to test. They help mathematicians posit conjectures to prove.⁵ And they help guide reasoning in other domains as well. For example, consider a prosecutor who attempts to convince a jury of a defendant's guilt by linking blood at the crime scene to the defendant. This prosecutor appeals to typicality. For typically, the presence of blood indicates the defendant's involvement in the crime. It would be quite atypical, though not impossible, for the defendant's blood to be present if the defendant were in no way involved.

²The notion of typicality has historical roots in the writings of Bernoulli (1713) and Cournot (1843). Both Bernoulli and Cournot formulated principles of typicality reasoning using the notion of probability, not typicality. Roughly, they argued that events with very high probability are 'morally certain', and events with very small probability are 'morally impossible' (Shafer & Vovk, 2006, p. 72). These principles of reasoning may be different from the analogous principles that replace the notion of probability with the notion of typicality, since arguably, probability and typicality are distinct (Goldstein, 2012; Wilhelm, forthcoming). But regardless, these principles of reasoning are at least direct ancestors of similar principles based on the notion of typicality.

³Every even integer greater than 2 can be expressed as the sum of two primes.

⁴Empirical verification of Goldbach's conjecture is catalogued at <http://sweet.ua.pt/tos/goldbach.html>.

⁵Pólya (1954) discussed something like this 'non-rigorous' side of mathematics in depth in his two-volume treatise *Mathematics and Plausible Reasoning*. Mazur (2012) has also explored the role of plausibility in mathematical practice.

Like typicality facts, typicality judgments can figure in our explanations: they can justify conjectural claims which have not yet been proven, but which there is sufficient evidence to support. Though Goldbach’s conjecture has remained unresolved for over three centuries, its plausibility is undisputed. There is lots of empirical evidence in its favor, though a formal proof remains elusive. Even without a proof, belief in the conjecture seems to be based on sound reasoning. But what is the logic of that reasoning?

Despite the ubiquity of typicality reasoning, no formal systems for the logic of typicality have been proposed.⁶ Most research on typicality either (i) proves results about what is typical – as in quantum mechanics and statistical mechanics – or (ii) explicates the notion of typicality – as in philosophy.⁷ There is comparatively little research on the logical principles which govern reasoning that relies upon typicality facts. There are no rigorous formal languages designed to model claims about what is and is not typical, or to assess the soundness of typicality judgments. There is no detailed formal semantic theory and no detailed proof theory for reasoning about typicality. Of course, there are rigorous formal systems for other sorts of reasoning. Propositional logic, and first-order logic, are formal systems for deductive reasoning. Bayesian theory is a formal system for probabilistic reasoning. But there are no analogous formalisms for typicality reasoning.⁸

There are at least two other reasons to develop a logical system for typicality reasoning. First, such a system would unite and systematize the different ways of quantifying typicality: it would show that the different typicality measures employed by Everett (1956/2012), Dürr et al. (1992), Reitmunn (2007), and others, are species of a common genus. In other words, a logical system for typicality reasoning would capture the formal unity of a wide variety of typicality results. It would reveal what many different approaches to typicality have in

⁶That is, there are no formal systems for our notion of typicality, according to which something is typical just in case, roughly, *nearly all* things of a certain sort are a certain way. There are proposed formal systems for other typicality notions. One notion of typicality, for instance, is given by the notion of ‘normal’: something is typical just in case it is normal (relative to the entities in some class). See (Booth et al., 2012) for a logic of this ‘normalcy’ notion of typicality.

⁷For recent work on the philosophical foundations of typicality, see (Frigg, 2011), (Frigg & Werndl, 2012), (Werndl, 2013), and (Wilhelm, forthcoming).

⁸Steps towards a formalism are taken by Goldstein et al. (2010, pp. 3217-3220).

common. Second, and relatedly, a logical system for typicality reasoning would formulate the basic logical principles that seem to govern the intuitive notion of typicality. Here is an example of one such principle: if $p \wedge q$ is typical then p is typical and q is typical, but if p is typical and q is typical then it does not follow that $p \wedge q$ is typical. Here is another: if either p is typical or q is typical then $p \vee q$ is typical, but if $p \vee q$ is typical then it does not follow that either p is typical or q is typical. A logical system for typicality reasoning would be a rigorous theory of principles which, like these two, govern all rational reasoning that relies on claims about what is typical.

So in this paper, we present two logical systems for typicality. The first is propositional: it supplements the standard language of propositional logic with a new sentence operator ‘*Typ*’. Intuitively, $Typ(p)$ says that p is typical. The second is type-theoretic: it supplements the standard language of intuitionistic Martin-Löf type theory (MLTT) (Martin-Löf, 1984) by introducing a new type former **Typ**. Intuitively, $\mathbf{Typ}(p)$ is a type corresponding to the proposition that p is typical, and each term of $\mathbf{Typ}(p)$ represents a justification for the judgment that p is typical.

In Section 2, we present the propositional formalism for typicality. We introduce the language, the semantics, and the proof theory for what we call *Typicality Propositional Logic* (TPL). We also establish some formal results. In Section 3, we present the type-theoretic formalism for typicality. We introduce the language of Martin-Löf type theory (MLTT), along with the semantics and proof theory for what we call *Typicality Intuitionistic Logic* (TIL).

2 Typicality Propositional Logic

In this section, we propose a formal logic for TPL. We introduce the language of that formalism—the basic vocabulary, and the well-formed formulas—in Section 2.1. In Section 2.2, we propose a semantic theory for this language, and we prove some simple yet illuminating

results. In Section 2.3, we propose a proof theory. In Section 2.4, we show that the proof theory is sound with respect to the semantic theory. Finally, in Section 2.5, we discuss some additional features of TPL.

2.1 The Language

TPL is the language of standard propositional logic supplemented with a typicality operator ‘*Typ*’. In particular, the logical vocabulary of TPL consists of three symbols: a binary sentence operator \rightarrow , a unary sentence operator \neg , and a unary sentence operator *Typ*. The non-logical vocabulary of TPL consists of infinitely many sentence letters and two bracket symbols: the sentence letters are p, q , and so on, and the bracket symbols are (and).

Well-formed formulas in the language of TPL are defined recursively, as follows.

1. Each sentence letter is a well-formed formula.
2. If ϕ is a well-formed formula, then $\neg\phi$ is a well-formed formula.
3. If ϕ is a well-formed formula, then *Typ*(ϕ) is a well-formed formula.
4. If ϕ and ψ are well-formed formulas, then $\phi \rightarrow \psi$ is a well-formed formula.
5. Nothing else is a well-formed formula.

It follows that all well-formed formulas in the language of propositional logic are well-formed formulas in the language of TPL.

2.2 The Semantics of TPL

The models of the well-formed formulas of TPL are called ‘TPL universes’. Each TPL universe is a pair $\langle \Gamma, \mathcal{V} \rangle$, where Γ is a large set and \mathcal{V} is a set of truth functions from well-formed formulas of TPL to $\{0, 1\}$. Intuitively, Γ is a set of possible states, or possible worlds, and \mathcal{V} is a set of functions which assign a truth value to each sentence, with 0 representing

‘false’ and 1 representing ‘true’. For each $w \in \Gamma$, there is exactly one function $f_w \in \mathcal{V}$. Intuitively, f_w expresses the facts about true propositions at w : for well-formed formula ϕ , for $w \in \Gamma$, and for $f_w \in \mathcal{V}$, $f_w(\phi) = 1$ says that ϕ is true at world w .

The rigorous definition of the truth functions proceeds in two steps. First, for each $w \in \Gamma$, let g_w be a truth function defined over all well-formed formulas of the language PL, where PL is the standard language for propositional logic. So each g_w is just a truth function of propositional logic.

Second, for each $w \in \Gamma$, extend g_w to a truth function f_w defined over all of TPL. The extension is defined in terms of double recursion: at each step in the recursion, in addition to defining the truth function f_w for that step, a set must also be defined. Roughly put, the defined set is the set of all elements of Γ at which a certain well-formed formula is true.

More precisely, let \mathcal{S} be a σ -algebra over Γ and let τ be a finite, non-zero measure over \mathcal{S} . Fix $\epsilon > 0$ such that $\epsilon \ll 1$. Then for each $w \in \Gamma$, the recursive definition of f_w is as follows.

- (1) For each well-formed formula ϕ in the language of PL,
 - (i) $f_w(\phi) = g_w(\phi)$, and
 - (ii) $\Gamma_\phi = \{w' \in \Gamma \mid f_{w'}(\phi) = 1\}$.
- (2) If ϕ is a well-formed formula in the language of TPL, then
 - (i) $f_w(\neg\phi) = 1$ if and only if $f_w(\phi) = 0$, and
 - (ii) $\Gamma_{\neg\phi} = \{w' \in \Gamma \mid f_{w'}(\phi) = 0\}$.
- (3) If ϕ and ψ are well-formed formulas in the language of TPL, then
 - (i) $f_w(\phi \rightarrow \psi) = 1$ if and only if $f_w(\phi) = 0$ or $f_w(\psi) = 1$, and
 - (ii) $\Gamma_{\phi \rightarrow \psi} = \{w' \in \Gamma \mid f_{w'}(\phi) = 0 \text{ or } f_{w'}(\psi) = 1\}$.
- (4) If ϕ is a well-formed formula in the language of TPL, then
 - (i) $f_w(Typ(\phi)) = 1$ if and only if there exists a set $X \in \mathcal{S}$ such that $X \subseteq \Gamma_\phi$ and $\frac{\tau(\Gamma \setminus X)}{\tau(\Gamma)} < \epsilon$, and
 - (ii) $\Gamma_{Typ(\phi)} = \{w' \in \Gamma \mid f_{w'}(Typ(\phi)) = 1\}$.

It follows from these clauses that for each well-formed formula ϕ , Γ_ϕ is the set of worlds in Γ

at which ϕ is true; that is, $\Gamma_\phi = \{w' \in \Gamma \mid f_{w'}(\phi) = 1\}$. For the purposes of the proof theory in Section 2.3, say that $Y \subseteq \Gamma$ is a ‘typical set’ if and only if there is an $X \in \mathcal{S}$ such that $X \subseteq Y$ and $\frac{\tau(\Gamma \setminus X)}{\tau(\Gamma)} < \epsilon$.

Here is an informal description of what these clauses say. According to the first clause, f_w agrees with g_w on the formulas of PL. The second clause uses the value of f_w on unnegated formulas in TPL to define the value of f_w on negated formulas in TPL. The third clause uses the value of f_w on pairs of formulas in TPL to define the value of f_w on conditionals created out of those formulas in TPL. The fourth clause is more involved: it uses facts about the measures of sets to define the value of f_w on typicality statements in TPL. Intuitively, the fourth clause says that $Typ(\phi)$ is true at a world if and only if the set of worlds at which ϕ holds – that is, the set Γ_ϕ – contains a ‘sufficiently large’ set X , where X is ‘sufficiently large’ if and only if X is measurable and the size of the set of elements not in X (divided by the size of the set Γ) is very small.

Note that in condition 4(i), the sizes of sets are quantified using a measure. But there are other ways of quantifying the sizes of sets. For example, a cardinality-theoretic version of 4(i) can be used for the case where Γ is infinite:⁹ $f_w(Typ(\phi)) = 1$ if and only if $|\Gamma \setminus \Gamma_\phi| < |\Gamma|$. In other words, ϕ is typical if and only if the set of $\neg\phi$ worlds has strictly smaller cardinality than the set of ϕ worlds.¹⁰ For a topological version of 4(i): $f_w(Typ(\phi)) = 1$ if and only if $\Gamma \setminus \Gamma_\phi$ is meager.¹¹ In other words, put roughly, ϕ is typical if and only if the set of $\neg\phi$ worlds is tightly packed, topologically, in Γ . And here is a version of 4(i) related to, but distinct from, the measure-theoretic version: given $\epsilon > 0$ such that $\epsilon \ll 1$, given a field of sets \mathcal{F} whose members are subsets of Γ , and given a finite non-zero finitely additive measure μ on \mathcal{F} , $f_w(Typ(\phi)) = 1$ if and only if there exists a set $X \in \mathcal{F}$ such that $X \subseteq \Gamma_\phi$ and $\frac{\mu(\Gamma \setminus X)}{\mu(\Gamma)} < \epsilon$.

⁹We stipulate that Γ is infinite because if Γ were finite, then this version of 4(i) would imply that $f_w(Typ(\phi)) = 1$ if and only if Γ_ϕ is nonempty. But this means that ϕ is typical if and only if ϕ is true in at least one world – and that is not how the intuitive notion of typicality works.

¹⁰As discussed in Wilhelm (forthcoming), this quantification of typicality is only suitable for modeling some typicality claims.

¹¹A set is meager if and only if it can be written as a countable union of nowhere dense sets. A set is nowhere dense if and only if its closure has empty interior.

According to condition 4(i), the truth of a typicality statement depends on the value of ϵ . The parameter ϵ serves as a measure of smallness: very roughly, ϕ is typical if and only if ϕ is false at a sufficiently small proportion of worlds, where a proportion of worlds is ‘sufficiently small’ if and only if that proportion is smaller than ϵ . So different choices of ϵ yield different truth conditions for claims about what is typical.

There are roughly two different views of the relationship between ϵ and the truth conditions for typicality statements. According to one view – call it the ‘context-dependent view’ of ϵ – many different values for ϵ are permissible, and there is a different version of condition 4(i) for each such value. The correct value for ϵ —the value, that is, which correctly captures the truth conditions of typicality statements—is determined by context. But different contexts can determine different values for ϵ . For example, consider a context in which physicists are studying the entropic behavior of gases. These physicists make veridical claims like “Typically, gases with such-and-such an initial macrostate evolve to a higher-entropy macrostate in thus-and-so amount of time”. In this context, the value of ϵ may be as low as 10^{-100} , because an absolutely massive number of gases exhibit the entropy-increasing behavior in question. In contrast, consider a context in which biologists are studying the behaviors of cells. These biologists make veridical claims like “Typically, cells with sodium-potassium pumps transport such-and-such many sodium ions in thus-and-so amount of time”. In this context, the value of ϵ may only be 10^{-4} , because the failure rate of the relevant transport processes is low but not as low as 10^{-100} .

According to another view – call it the ‘context-independent view’ of ϵ – only one value of ϵ is correct: only one version of condition 4(i) can be used in typicality reasoning. That value for ϵ is always the same; it does not vary with context. For example, the value of ϵ in the context of physicists studying gases is the same as the value of ϵ in the context of biologists studying cells.

We prefer the context-dependent view of ϵ . Typicality, on our view, is a context-dependent notion: whether or not something is typical varies with the standards of the

context in question. A formal system for typicality should capture that contextual variability. This does not make TPL any less exact, however; this does not make TPL imprecise. By invoking the parameter ϵ , TPL exactly and precisely captures the inexactness and imprecision which is inherent in the notion of typicality. A formal system which fails to capture the contextual variability of typicality is not any more exact or precise than TPL. Such a formal system is, in fact, not a formal system for typicality at all.

But in this paper, for the sake of brevity, we will not explore different ways of varying the parameter ϵ . So in the theorems to follow, we will assume that ϵ has been fixed to some particular value. We hope that future work will explore the consequences of allowing the value of ϵ to vary.¹²

We now present the account of truth at a world in a TPL universe, an account of truth in a TPL universe, and an account of logical truth. Let $M_0 = \langle \Gamma_0, \mathcal{V}_0 \rangle$ be a TPL universe, let $w_0 \in \Gamma_0$ be a world, and let ϕ be a well-formed formula in the language of TPL. First, ϕ is ‘true in M_0 at w_0 ’ if and only if $f_{w_0}(\phi) = 1$. In symbols: $\models_{M_0, w_0} \phi$. Second, ϕ is ‘true in M_0 ’ if and only if for every $w \in \Gamma_0$, ϕ is true in M_0 at w . In symbols: $\models_{M_0} \phi$. Third, ϕ is ‘logically true’ if and only if for every TPL universe M , ϕ is true in M . In symbols: $\models \phi$.

Define logical entailment in TPL as follows. Let Σ be a set of well-formed formulas in the language of TPL. Let M be a TPL universe. Say that ‘ Σ logically entails ψ in M ’ if and only if the following holds: if ϕ is true in M for each $\phi \in \Sigma$, then ψ is true in M . In symbols: $\Sigma \models_M \psi$. And say that ‘ Σ logically entails ψ ’ if and only if for each TPL universe M , Σ logically entails ψ in M . In symbols: $\Sigma \models \psi$.¹³

¹²See footnote 13 for one example of how the present theory can be adapted to allow for contextually variable ϵ .

¹³This definition of logical entailment depends upon a specific choice of $\epsilon > 0$ such that $\epsilon \ll 1$, since ϵ was fixed to a single value for the purposes of the recursive definitions of the functions f_w that assign truth values to sentences at worlds in TPL universes. But this definition of logical entailment can easily be adapted to allow for multiple values of ϵ . To do so, simply call this definition—the one which depends upon a fixed ϵ —an ‘ ϵ -relative’ definition of logical entailment. Then say that the set of sentences Σ logically entails the sentence ψ if and only if for all $\epsilon > 0$ such that $\epsilon < \frac{1}{2}$, Σ logically entails ψ according to the ϵ -relative definition of logical entailment. The choice of $\frac{1}{2}$ as an upper bound is not forced: a different number, for instance $\frac{1}{10^4}$ or $\frac{1}{10^{100}}$, could be used instead. But $\frac{1}{2}$ seems like the least arbitrary choice for this non- ϵ -relative definition of logical entailment. For any other choice, it seems reasonable to ask why ϵ cannot be just a little bit higher; and it is not obvious what the answer would be. But clearly, ϵ cannot be $\frac{1}{2}$ or greater: for if something is

A few simple results will help to clarify the nature of typicality in this semantic theory. The following lemma shows that typicality statements are ‘all or nothing’: either $Typ(\phi)$ is true at each world in the universe at issue or $Typ(\phi)$ is false at each world in the universe at issue.

Lemma 1. *Let $\langle \Gamma, \mathcal{V} \rangle$ be a TPL universe. For any well-formed formula ϕ in the language of TPL, either $\Gamma_{Typ(\phi)} = \Gamma$ or $\Gamma_{Typ(\phi)} = \emptyset$.*

Proof. Take any well-formed formula ϕ in the language of TPL. Suppose $\Gamma_{Typ(\phi)} \neq \Gamma$. Then there is a world $w \in \Gamma$ such that $f_w(Typ(\phi)) = 0$. Recall that for any $w' \in \Gamma$, $f_{w'}(Typ(\phi)) = 1$ if and only if there exists a set $X \in \mathcal{S}$ such that $X \subseteq \Gamma_\phi$ and $\frac{\tau(\Gamma \setminus X)}{\tau(\Gamma)} < \epsilon$. So since $f_w(Typ(\phi)) = 0$, there is no set X which satisfies those conditions. Therefore, for every $w' \in \Gamma$, $f_{w'}(Typ(\phi)) = 0$. And therefore, $\Gamma_{Typ(\phi)} = \emptyset$. So either $\Gamma_{Typ(\phi)} = \Gamma$ or $\Gamma_{Typ(\phi)} = \emptyset$. \square

Lemma 1 justifies the use of the phrase ‘typicality facts’ to refer to typicality statements in TPL. When a typicality statement holds at some world, it holds at all worlds. If a typicality statement fails to hold at a world, it fails to hold at all worlds. Therefore, the truth value of a typicality statement is constant over all worlds in a TPL universe, and is thus a *fact* about that universe.

The following theorem shows that iterated typicality claims do not change in truth value: that is, $Typ(Typ(\phi))$ holds if and only if $Typ(\phi)$ holds.

Theorem 1. *For any well-formed formula ϕ in the language of TPL, for any TPL universe $M = \langle \Gamma, \mathcal{V} \rangle$, and for any $w \in \Gamma$, $Typ(Typ(\phi))$ is true in M at w if and only if $Typ(\phi)$ is true in M at w .*

Proof. By definition, $Typ(Typ(\phi))$ is true in M at w if and only if $f_w(Typ(Typ(\phi))) = 1$. Similarly, $Typ(\phi)$ is true in M at w if and only if $f_w(Typ(\phi)) = 1$. So to establish the theorem, it suffices to show that $f_w(Typ(Typ(\phi))) = f_w(Typ(\phi))$.

typical, then at the very least, it must be true more than half of the time.

By Lemma 1, either $\Gamma_{Typ(\phi)} = \emptyset$ or $\Gamma_{Typ(\phi)} = \Gamma$. To start, suppose that $\Gamma_{Typ(\phi)} = \emptyset$. Then for each set $X \in \mathcal{S}$ such that $X \subseteq \Gamma_{Typ(\phi)}$, $X = \emptyset$. Therefore, for such an X ,

$$\begin{aligned} \frac{\tau(\Gamma \setminus X)}{\tau(\Gamma)} &= \frac{\tau(\Gamma \setminus \emptyset)}{\tau(\Gamma)} \\ &= \frac{\tau(\Gamma)}{\tau(\Gamma)} \\ &= 1 \\ &\not\prec \epsilon. \end{aligned}$$

And so by definition, $f_w(Typ(Typ(\phi))) = 0$. In addition, if $\Gamma_{Typ(\phi)} = \emptyset$, then $f_w(Typ(\phi)) = 0$ by definition. Therefore, if $\Gamma_{Typ(\phi)} = \emptyset$, then $f_w(Typ(Typ(\phi))) = f_w(Typ(\phi))$.

Now suppose that $\Gamma_{Typ(\phi)} = \Gamma$. Then there is a set $X \in \mathcal{S}$ such that $X \subseteq \Gamma_{Typ(\phi)}$ and $\frac{\tau(\Gamma \setminus X)}{\tau(\Gamma)} < \epsilon$; namely, the set $X = \Gamma$. So by definition, $f_w(Typ(Typ(\phi))) = 1$. In addition, if $\Gamma_{Typ(\phi)} = \Gamma$, then $f_w(Typ(\phi)) = 1$ by definition. So if $\Gamma_{Typ(\phi)} = \Gamma$, then $f_w(Typ(Typ(\phi))) = f_w(Typ(\phi))$.

Therefore, regardless of what $\Gamma_{Typ(\phi)}$ is, $f_w(Typ(Typ(\phi))) = f_w(Typ(\phi))$.

□

The following corollary establishes the corresponding result for the sets $\Gamma_{Typ(Typ(\phi))}$ and $\Gamma_{Typ(\phi)}$.

Corollary 1. *For any well-formed formula ϕ in the language of TPL, and for any TPL universe $M = \langle \Gamma, \mathcal{V} \rangle$, $\Gamma_{Typ(Typ(\phi))} = \Gamma_{Typ(\phi)}$.*

Proof. By definition, $\Gamma_{Typ(Typ(\phi))} = \{w' \in \Gamma \mid f_{w'}(Typ(Typ(\phi))) = 1\}$ and $\Gamma_{Typ(\phi)} = \{w' \in \Gamma \mid f_{w'}(Typ(\phi)) = 1\}$. Theorem 1 implies that for each $w \in \Gamma$, $f_w(Typ(Typ(\phi))) = f_w(Typ(\phi))$.

Therefore, $\Gamma_{Typ(Typ(\phi))} = \Gamma_{Typ(\phi)}$.

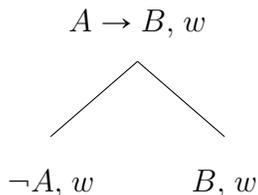
□

2.3 Tableaus for TPL

In this subsection, we propose a proof theory for TPL based on tableaus; call it a ‘TPL proof theory’.¹⁴ The method consists of a series of rules for decomposing well-formed formulas in the language of TPL. The decomposition resembles a tree-like structure; because of that, tableau proofs are often called ‘proof trees’.

Each line in a proof tree consists of two parts: a formula and an index. The formula is just a well-formed formula in the language of TPL. The index is a syntactic parameter that keeps track of certain information about that well-formed formula. Intuitively, the index specifies the worlds at which the formula is true. There are three different types of indices: the ‘all’ index, denoted by the symbol a ; ‘nearly all’ indices, denoted by symbols like n , n' , and so on; and ‘world’ indices, denoted by symbols like w , w' , and so on. Intuitively, a line of the form ‘ A, a ’ says that A is true at all worlds. A line of the form ‘ A, n ’ says that there is a typical set such that A is true at each world in that set: think of n as that typical set’s name. And a line of the form ‘ A, w ’ says that A is true at world w .

We now present the decomposition rules for the proof trees. In each rule, A and B are schematic letters for well-formed formulas of TPL. The first rule is the ‘conditional’ rule:



Intuitively, interpret this rule as saying that if $A \rightarrow B$ is true at world w , then either $\neg A$ is true at w or B is true at w . This rule is, however, purely syntactic. Later we will justify this interpretation of the syntax. Doing so is the first step towards showing that this TPL proof theory is sound with respect to the TPL semantic theory discussed in Section 2.2.

¹⁴For discussion of tableau methods in propositional logic, first-order logic, and modal logic, see Smullyan (1968) and Priest (2008).

The next three rules are ‘negated conditional’ rules:

$$\begin{array}{ccc}
 \neg(A \rightarrow B), a & \neg(A \rightarrow B), n & \neg(A \rightarrow B), w \\
 | & | & | \\
 A, a & A, n & A, w \\
 \neg B, a & \neg B, n & \neg B, w
 \end{array}$$

Intuitively, interpret the rule with index ‘ a ’ as saying the following: if $\neg(A \rightarrow B)$ is true at every world, then A and $\neg B$ are true at every world. Interpret the rule with index ‘ n ’ as saying the following: if $\neg(A \rightarrow B)$ is true at every world in the typical set named by n , then A and $\neg B$ are true at every world in n . And interpret the rule with index ‘ w ’ as saying the following: if $\neg(A \rightarrow B)$ is true at world w , then A and $\neg B$ are true at w .

The next three rules are ‘negation’ rules:

$$\begin{array}{ccc}
 \neg\neg A, a & \neg\neg A, n & \neg\neg A, w \\
 | & | & | \\
 A, a & A, n & A, w
 \end{array}$$

Intuitively, interpret the rule with index ‘ a ’ as saying the following: if $\neg\neg A$ is true at every world, then A is true at every world. Interpret the rule with index ‘ n ’ as saying the following: if $\neg\neg A$ is true at every world in the typical set named by n , then A is true at every world in n . And interpret the rule with index ‘ w ’ as saying the following: if $\neg\neg A$ is true at world w , then A is true at w .

The next three rules are ‘typicality’ rules:

$$\begin{array}{ccc}
 Typ(A), a & Typ(A), n & Typ(A), w \\
 | & | & | \\
 A, n' & A, n' & A, n'
 \end{array}$$

where in each rule, n' is a new ‘nearly all’ index, one not used earlier in the tableau. Think of n' as naming the set of worlds at which A holds. So intuitively, the first rule says that if A is typical at each world, then there is a typical set n' such that A holds at each world in n' . The second and third rules say something similar, except that they start with slightly different assumptions: the second starts with the assumption that A is typical at all worlds in some typical set n , and the third starts with the assumption that A is typical at some world w .

We require that n' be a *new* ‘nearly all’ index—we require, in other words, that n' not show up earlier in the proof tree—for the same reasons that in the tableau method for first-order logic, one always uses a new constant symbol ‘ c ’ when decomposing a formula of the form $\exists x\phi(x)$ into $\phi(c)$.¹⁵ $Typ(A)$ is, in a sense, an implicitly existential sort of formula: $Typ(A)$ holds in a TPL universe $\langle \Gamma, \mathcal{V} \rangle$ if and only if *there exists* a typical set $X \subseteq \Gamma$ such that A is true at each world in X . So as with other such existential rules in other tableau methods, we require that the name for the item whose existence is asserted in the decomposition—the name ‘ n' ’, in this case, which denotes a typical set of worlds such that A is true at each world in that set—not be associated with any other properties.¹⁶

The next three rules are ‘negated typicality’ rules:

$$\begin{array}{ccc}
 \neg Typ(A), a & \neg Typ(A), n & \neg Typ(A), w \\
 \left| & \left| & \left| \\
 w' \in n' & w' \in n' & w' \in n' \\
 \neg A, w' & \neg A, w' & \neg A, w'
 \end{array}$$

where n' is a ‘nearly all’ index, and w' is a ‘world’ index which was not invoked earlier in

¹⁵For a nice discussion, see Smullyan (1968, p. 54).

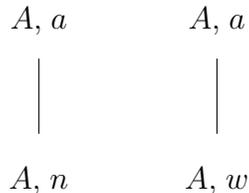
¹⁶If we did not require this, then the tableau method would not match the semantics of Section 2.2 in the right way. Proof trees would say more than is licensed by the semantics for TPL. For instance, suppose that some other typical set n^* is mentioned earlier in the proof tree, and B is true at each n^* . We do not yet know whether or not A is also true at each world in n^* . So if we decomposed the line ‘ $Typ(A), a$ ’ as A, n^* , we would be saying something stronger than, at this point in the tree, we legitimately can. We would be saying, for example, that A and B are both true at each world in n^* . But we do not know that. We only know that A is typical and that B is typical: we do not know that their conjunction is typical (indeed, it may not be).

the proof. So intuitively, the first rule says that if it is not the case that A is typical at each world, then for any typical set n' there is a world in n' such that A is false at that world. The second and third rules say something similar, except that they start with slightly different assumptions: the second starts with the assumption that it is not the case that A is typical at all worlds in some typical set n , and the third starts with the assumption that it is not the case that A is typical in some world w .

Note that in this rule, the index n' may or may not have appeared earlier on. In fact, for various pragmatic reasons, this rule should often be applied whenever a new ‘nearly all’ index appears in the proof tree. For instance, if the tree features the three lines ‘ B, n_1 ’, ‘ C, n_2 ’, and ‘ $\neg Typ(A), a$ ’, then apply the first negated typicality rule twice: once for the case where $n' = n_1$, and once for the case where $n' = n_2$.¹⁷

In addition, note that the index w' must *not* have shown up earlier in the proof tree. The reason is that $\neg Typ(A)$ contains an implicit existential: $\neg Typ(A)$ holds in a TPL universe $\langle \Gamma, \mathcal{V} \rangle$ if and only if for each typical set $X \subseteq \Gamma$, *there exists* a world $w' \in X$ such that A is false at w' .¹⁸ So as with other such existential rules in other tableau methods, we require that the name for the item whose existence is asserted in the decomposition—the name ‘ w' ’, in this case, which denotes a world in Γ —not be associated with any other properties.

The remaining rules concern indices in particular. The first two are ‘all-to-less’ index rules:



¹⁷Whether or not this rule should *always* be applied for *each* new ‘nearly all’ index will depend, ultimately, on pragmatic considerations. Some tableau proofs, for some logical systems, can go on forever: see the discussion by Smullyan (1968). So for pragmatic reasons, this rule should not be applied in a way that, in conjunction with the other rules, generates an infinite tree – if at all possible.

¹⁸For if there were a typical set X in which no such world existed – that is, where A was true at each world in X – then $Typ(A)$ would be true in the TPL universe in question.

where n is a ‘nearly all’ index and w is a ‘world’ index. Intuitively, the first all-to-less index rule says that if A holds at every world, then for any typical set n , A holds at each world in n . The second all-to-less index rule says that if A holds at every world, then for any particular world w , A holds at w .

The next rule is an ‘instantiation’ index rule:

$$\begin{array}{c} A, n \\ w \in n \\ | \\ A, w \end{array}$$

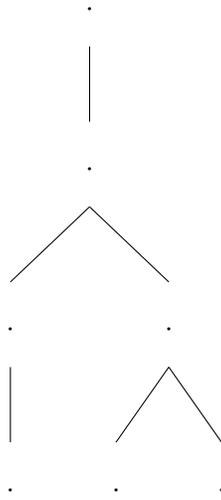
Intuitively, the instantiation index rule says that if A is true at each world in the typical set n , and if w is in n , then A is true at w .

The final rule is the ‘world introduction’ rule:

$$\begin{array}{c} | \\ w \in n \end{array}$$

where n is a ‘nearly all’ index which may or may not have been invoked earlier in the proof, and w is a ‘world’ index which was not invoked earlier in the proof. Intuitively, the world introduction rule says that for any given typical set n , there is a world w in n .

It will be helpful to have some terminology for various features of proof trees. In general, proof trees have structures like the following:



The dots in the above image, which in TPL proof trees consist of a collection of pairs of formulas and indices, are called ‘nodes’. The ‘initial’ node is the one at the tree’s top. A vertical line between two nodes indicates that one of the decomposition rules has been applied to the node at the top, yielding the node at the bottom. A ‘branch’ of a proof tree is a path from the initial node to a lower node. A ‘closed’ branch is a branch that contains two contradictory formulas with the same index: for instance, a branch that contains ‘ A, w ’ and ‘ $\neg A, w$ ’ is closed. An ‘open’ branch is a branch which is not closed. A ‘closed’ tree is a proof tree such that every branch is closed. An ‘open’ tree is a proof tree which is not closed.

Let us now define deductive TPL proofs. Let $\psi, \phi_1, \phi_2, \dots, \phi_n$ be well-formed formulas in the language of TPL. Say that ψ is deducible from $\phi_1, \phi_2, \dots, \phi_n$ if and only if there is a proof tree \mathfrak{p} which satisfies the following conditions.

1. The initial node of \mathfrak{p} consists of the following collection of pairs of formulas and indices:

$$\begin{array}{c} \phi_1, a \\ \phi_2, a \\ \vdots \\ \phi_n, a \end{array}$$

$\neg\psi, w$

2. The tree \mathfrak{p} is closed.

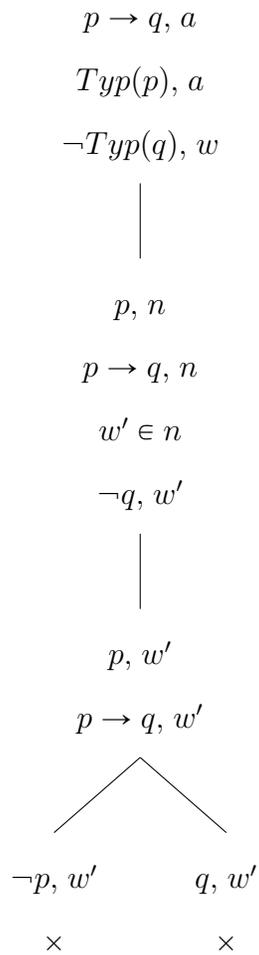
Let $\phi_1, \phi_2, \dots, \phi_n \vdash \psi$ denote that ψ is deducible from $\phi_1, \phi_2, \dots, \phi_n$. More generally, if Σ is a finite set of well-formed formulas, let $\Sigma \vdash \psi$ denote that ψ is deducible from the formulas in Σ . If ϕ is a well-formed formula in the language of TPL, then ϕ is a ‘theorem’ if and only if $\vdash \phi$; that is, ϕ is a theorem if and only if ϕ can be deduced without invoking any assumptions.

As an example, let us prove that $p \vdash Typ(p)$. Here is the proof tree. The \times at the bottom indicates that the branch is closed.

$$\begin{array}{c} p, a \\ \neg Typ(p), w \\ | \\ w' \in n \\ \neg p, w' \\ | \\ p, w' \\ \times \end{array}$$

This tree was obtained by first applying the negated typicality rule for a world index to the second line, and then applying the all-to-less rule for a world index to the first line. The proof tree shows that if p is true at every world, and if $Typ(p)$ is false at some world, then a contradiction obtains. And intuitively, that is correct. For if something is true at every world, then that something must be typical at every world.

As another example, let us prove that $p \rightarrow q, Typ(p) \vdash Typ(q)$. Here is the proof tree.



This tree shows that if $p \rightarrow q$ is true at every world, if $Typ(p)$ is true at every world, and if $Typ(q)$ is false at some world, then a contradiction is reached. That is, if $p \rightarrow q$ is true at every world, and if $Typ(p)$ is true at every world, then $Typ(q)$ must also be true at every world.

As a final example, let us see why $\not\vdash p \rightarrow Typ(p)$. Here is the relevant tree.

$$\begin{array}{c}
\neg(p \rightarrow Typ(p)), w \\
| \\
p, w \\
\neg Typ(p), w \\
| \\
w' \in n \\
\neg p, w'
\end{array}$$

The tree is not closed. Of course, more rules could still be applied. We could apply the negated typicality rule for a world index again, or we could apply the world introduction rule. But clearly, repeated applications of these rules will never close the tree. So there is no proof of $p \rightarrow Typ(p)$. And that is intuitively the right result, since the formula ' $p \rightarrow Typ(p)$ ' need not be true at every world in every TPL universe. Consider a TPL universe where p is false at all worlds but one. Then $p \rightarrow Typ(p)$ is false at that world: for at that world, p is true and $Typ(p)$ is false. Therefore, $p \rightarrow Typ(p)$ is not true in this TPL universe. And therefore, $p \rightarrow Typ(p)$ is not valid.

2.4 Soundness

When explaining the intuitions motivating the decomposition rules, we often invoked semantic notions. In this section, we precisify the relationship between the proof theory of Section 2.3 and the semantic theory of Section 2.2. In particular, we show that the proof theory is sound with respect to the semantic theory.

To start, consider the following definition.

Definition 1 (Faithfulness). *Let $M = \langle \Gamma, \mathcal{V} \rangle$ be a TPL universe, and let b be any branch of a tableau. Say that M is 'faithful' to b if and only if there is a function f which takes each*

world index w on b to a world $f(w) \in \Gamma$, and takes each nearly all index n on b to a typical set $f(n) \subseteq \Gamma$, such that the following conditions hold.

1. For every node of the form A, a on b , A is true at each world in Γ .
2. For every node of the form A, n on b , A is true at each world in $f(n)$.
3. For every node of the form A, w on b , A is true at $f(w)$.
4. For every node of the form $w \in n$ on b , $f(w) \in f(n)$.

The following lemma will help to prove the soundness of the tableau method. It uses the notion of faithfulness to show that the decomposition rules of the proof theory preserve truth.

Lemma 2. *Let b be a branch of a tableau, and let $M = \langle \Gamma, \mathcal{V} \rangle$ be a TPL universe. If M is faithful to b , and a tableau rule is applied to b , then this tableau rule yields at least one extended branch to which M is faithful.*

Proof. Let f be the function which witnesses the fact that M is faithful to b . The proof of this lemma proceeds by checking every case: we show that for each tableau rule, an application of that tableau rule yields a branch b' to which M is faithful.

Suppose we apply the conditional rule for $A \rightarrow B, w$ to b . Since $A \rightarrow B, w$ is on b , $A \rightarrow B$ is true at $f(w)$. So either $\neg A$ is true at $f(w)$ or B is true at $f(w)$. So for at least one of these extensions of b —the one that extends to the left in the tree diagram, which has line $\neg A, w$; or the one that extends to the right in the tree diagram, which has line B, w — M is faithful to that extension.

Suppose we apply the negated conditional rule for $\neg(A \rightarrow B), a$ to b : the resulting branch b' has all the lines of b as well as the lines A, a and $\neg B, a$. Since $\neg(A \rightarrow B), a$ is on b and M is faithful to b , $\neg(A \rightarrow B)$ is true at each world in Γ . So A is true at each world in Γ and $\neg B$ is true at each world in Γ . Therefore, M is faithful to b' . Similarly, suppose we apply the negated conditional rule for $\neg(A \rightarrow B), n$ to b : the resulting branch b' has all the lines of b as well as the lines A, n and $\neg B, n$. Since $\neg(A \rightarrow B), n$ is on b and M is faithful to b , $\neg(A \rightarrow B)$ is true at each world in $f(n)$. So A is true at each world in $f(n)$ and $\neg B$ is true

at each world in $f(n)$. Therefore, M is faithful to b' . Finally, suppose we apply the negated conditional rule for $\neg(A \rightarrow B), w$ to b : the resulting branch b' has all the lines of b as well as the lines A, w and $\neg B, w$. Since $\neg(A \rightarrow B), w$ is on b and M is faithful to b , $\neg(A \rightarrow B)$ is true at $f(w)$. So A is true at $f(w)$ and $\neg B$ is true at $f(w)$. Therefore, M is faithful to b' .

Suppose we apply the negation rule for $\neg\neg A, a$ to b : the resulting branch b' has all the lines of b as well as the line A, a . Since $\neg\neg A, a$ is on b and M is faithful to b , $\neg\neg A$ is true at each world in Γ . So A is true at each world in Γ . Therefore, M is faithful to b' . Similarly, suppose we apply the rule for $\neg\neg A, n$ to b : the resulting branch b' has all the lines of b as well as the line A, n . Since $\neg\neg A, n$ is on b and M is faithful to b , $\neg\neg A$ is true at each world in $f(n)$. So A is true at each world in $f(n)$. Therefore, M is faithful to b' . Finally, suppose we apply the rule for $\neg\neg A, w$ to b : the resulting branch b' has all the lines of b as well as the line A, w . Since $\neg\neg A, w$ is on b and M is faithful to b , $\neg\neg A, w$ is true at $f(w)$. So A is true at $f(w)$. Therefore, M is faithful to b' .

Suppose we apply the typicality rule for $Typ(A), a$ to b : the resulting branch b' has all the lines of b as well as the line A, n' , where n' is a ‘nearly all’ index which does not appear in b . Since $Typ(A), a$ is on b , $Typ(A)$ is true at each world in Γ . So there is a typical set $X \subseteq \Gamma$ such that A is true at each world in X . Extend f to a function f' as follows: for all w in the domain of f , $f'(w) = f(w)$; for all n in the domain of f , $f'(n) = f(n)$; and $f'(n') = X$.¹⁹ By definition 1, M is faithful to b' : that faithfulness is witnessed by f' .

Now suppose we apply the typicality rule for $Typ(A), n$ to b : the resulting branch b' has all the lines of b as well as the line A, n' , where n' is a ‘nearly all’ index which does not appear in b . Since $Typ(A), n$ is on b and M is faithful to b , $Typ(A)$ is true at each world in some typical set $Y \subseteq \Gamma$. Lemma 1, in conjunction with the fact that Y is nonempty, implies that $Typ(A)$ is true at each world in Γ . So by the exact same reasoning as before, M is faithful to b' .

¹⁹This last clause uses the fact that n' is a new ‘nearly all’ index, one which does not appear earlier on b' ; that is, n' is a ‘nearly all’ index which does not show up on b . For if n' did show up on b , then n' would already be in the domain of f , so we would have to set $f'(n') = f(n')$. That is, we would not be free to set $f'(n') = X$.

Similarly, suppose we apply the typicality rule for $Typ(A), w$ to b : the resulting branch b' has all the lines of b as well as the line A, n' , where n' is a ‘nearly all’ index which does not appear in b . Since $Typ(A), w$ is on b and M is faithful to b , $Typ(A)$ is true at world w . So once again, Lemma 1 implies that $Typ(A)$ is true at each world in Γ . And so by the exact same reasoning as before, M is faithful to b' .

Suppose we apply the negated typicality rule for $\neg Typ(A), a$ to b : the resulting branch b' has all the lines of b as well as the two lines $w' \in n'$ and $\neg A, w'$, where w' is a ‘world’ index which does not show up on b and n' is a ‘nearly all’ index which may or may not show up on b . Suppose n' shows up earlier on b . Then $f(n')$ is already defined. Since $\neg Typ(A), a$ is on b , A is false at some world q in $f(n')$. For if A were true at each world in $f(n')$, then by condition 4(i) of the recursive definition of truth functions in Section 2.2, and by Lemma 1, $Typ(A)$ would be true at each world in Γ . Extend f to a function f' as follows: for all n in the domain of f , $f'(n) = f(n)$; for all w in the domain of f , $f'(w) = f(w)$; and $f'(w') = q$. Then M is faithful to b' , as witnessed by f' . Now suppose that n' does not show up earlier on b . Since $\neg Typ(A), a$ is on b , Lemma 1 implies that $Typ(A)$ is false at each world in Γ . So take any typical set $X \subseteq \Gamma$. As before, since $Typ(A)$ is false at each world in Γ , there must be a world q in X at which A is false. Extend f to a function f' as follows: for all n in the domain of f , $f'(n) = f(n)$; for all w in the domain of f , $f'(w) = f(w)$; $f'(n') = X$; and $f'(w') = q$. Then M is faithful to b' , as witnessed by f' .

Because of lemma 1, the same line of argument works for the negated typicality rules for $\neg Typ(A), n$ and $\neg Typ(A), w$. So we do not present those arguments here. In each case, the conclusion is that M is faithful to the extension of b .

Suppose we apply the first all-to-less index rule for A, a to b : the resulting branch b' has all the lines of b as well as the line A, n . If n shows up on b , then $f(n)$ is already defined; $f(n) = Y$, say, where $Y \subseteq \Gamma$ is a typical set. Since A, a is on b and M is faithful to b , A is true at each world in Γ . So A is true at each world in Y , and therefore, M is faithful to b' . If n does not show up on b , then take any typical set $X \subseteq \Gamma$. Extend f to a function f' as

follows: for all n in the domain of f , $f'(n) = f(n)$; for all w in the domain of f , $f'(w) = f(w)$; and $f'(n') = X$. By definition 1, M is faithful to b' : that faithfulness is witnessed by f' .

Similarly, suppose we apply the second all-to-less rule for A, a to b : the resulting branch b' has all the lines of b as well as the line A, w . If w shows up earlier on b , then $f(w)$ is already defined. So $f(w) = q$, say, where $q \in \Gamma$. And since A, a is on b , A is true at each world in Γ . So A is true at q , and therefore, M is faithful to b' . If w does not show up earlier on b , then pick a world $q \in \Gamma$. Extend f to a function f' as follows: for all n in the domain of f , $f'(n) = f(n)$; for all w' in the domain of f , $f'(w') = f(w')$; and $f'(w) = q$. By definition 1, M is faithful to b' : that faithfulness is witnessed by f' .

Suppose we apply the instantiation rule for A, n and $w \in n$ to b : the resulting branch b' has all the lines of b as well as the line A, w . Since A, n is true on b and M is faithful to b , A is true at each world in n . Since f witnesses the fact that M is faithful to b , $f(w) \in f(n)$. Therefore, A is true at $f(w)$. So M is faithful to b' .

Finally, suppose we apply the world introduction rule to b : the resulting branch b' has all the lines of b as well as the line $w \in n$, where n is a ‘nearly all’ index that may or may not appear earlier on b and w is a ‘world’ index which does not show up earlier on b . Suppose that n appears earlier on b . Then note that if $f(n)$ is empty, then $\frac{\tau(\Gamma \setminus f(n))}{\tau(\Gamma)} = 1 \not\leq \epsilon$, contradiction. Thus, $f(n)$ is nonempty. So extend f to f' as follows: for all n^* in the domain of f , $f'(n^*) = f(n^*)$; for all w' in the domain of f , $f'(w') = f(w')$; and $f'(w)$ is some chosen element of $f(n)$. Therefore, M is faithful to b' : that faithfulness is witnessed by f' . Now suppose that n does not appear earlier on b . Then pick a typical set $X \subseteq \Gamma$, and extend f to f' as follows: for all n^* in the domain of f , $f'(n^*) = f(n^*)$; for all w' in the domain of f , $f'(w') = f(w')$; $f'(n) = X$; and $f'(w)$ is some chosen element of $f'(n)$.

□

Now for the proof of soundness.

Theorem 2 (Soundness). *If $\phi_1, \phi_2, \dots, \phi_n \vdash \psi$, then $\phi_1, \phi_2, \dots, \phi_n \models \psi$.*

Proof. The proof establishes the contrapositive. So suppose that $\phi_1, \phi_2, \dots, \phi_n \not\models \psi$. Then there is a TPL universe $M = \langle \Gamma, \mathcal{V} \rangle$ such that $\models_M \phi_1, \models_M \phi_2, \dots, \models_M \phi_n$, but for some world $w \in \Gamma$, $\models_{M,w} \neg\psi$. Consider a tableau with $\phi_1, \phi_2, \dots, \phi_n, \neg\psi$ at its start. Since M models $\phi_1, \phi_2, \dots, \phi_n, \neg\psi$, M is faithful to this initial branch: the witnessing function f maps the world index for $\neg\psi$ in the initial node of this tableau to the world w . By lemma 2, every application of a tableau rule to this initial branch yields at least one extended branch to which M is faithful. So there is no finite sequence of applications of tableau rules to this initial branch which closes the whole tree. So $\phi_1, \phi_2, \dots, \phi_n \not\models \psi$. □

2.5 Additional Results

In this subsection, we present a few more facts about TPL. In Section 2.5.1, we establish some simple results about the typicality of conjunctions and disjunctions. In Section 2.5.2, we establish two partial deduction theorems, one syntactic and one semantic. In Section 2.5.3, we outline an important difference between TPL and standard systems of modal logic.

2.5.1 Conjunction and Disjunction

In this subsection, we present some results concerning typicality statements about conjunction and disjunction. As usual, define $A \wedge B$ as $\neg(A \rightarrow \neg B)$, and define $A \vee B$ and $\neg A \rightarrow B$. It is straightforward to show that for any TPL universe $\langle \Gamma, \mathcal{V} \rangle$, any world $w \in \Gamma$, and any well-formed formulas A and B , $f_w(A \wedge B) = 1$ if and only if $f_w(A) = 1$ and $f_w(B) = 1$, and $f_w(A \vee B) = 1$ if and only if $f_w(A) = 1$ or $f_w(B) = 1$.

The typicality results about conjunction and disjunction are as follows.

Theorem 3. *Let ϕ and ψ be well-formed formulas of TPL. Then the following two conditions hold.*

1. $Typ(\phi \wedge \psi) \models Typ(\phi) \wedge Typ(\psi)$.

2. $Typ(\phi) \vee Typ(\psi) \models Typ(\phi \vee \psi)$.

Proof. To begin, let us establish condition 1. Let $M = \langle \Gamma, \mathcal{V} \rangle$ be a TPL universe, and suppose that $Typ(\phi \wedge \psi)$ is true in M . Take any $w \in \Gamma$. Since $Typ(\phi \wedge \psi)$ is true in M , $Typ(\phi \wedge \psi)$ is true at w . So by the recursive definition of truth in Section 2.2, there is a set $X \in \mathcal{S}$ such that $X \subseteq \Gamma_{\phi \wedge \psi}$ and $\frac{\tau(\Gamma \setminus X)}{\tau(\Gamma)} < \epsilon$. For each $w' \in \Gamma_{\phi \wedge \psi}$, $w' \in \Gamma_{\phi}$ and $w' \in \Gamma_{\psi}$, since if $\phi \wedge \psi$ is true at w' then ϕ is true at w' and ψ is true at w' . Therefore, $\Gamma_{\phi \wedge \psi} \subseteq \Gamma_{\phi}$ and $\Gamma_{\phi \wedge \psi} \subseteq \Gamma_{\psi}$. It follows that $X \subseteq \Gamma_{\phi}$ and $X \subseteq \Gamma_{\psi}$. Therefore, by the recursive definition of truth in Section 2.2, $Typ(\phi)$ is true at w and $Typ(\psi)$ is true at w . So Lemma 1 implies that both $Typ(\phi)$ and $Typ(\psi)$ are true at each world in M . Therefore, $Typ(\phi) \wedge Typ(\psi)$ is true in M . And therefore, $Typ(\phi \wedge \psi) \models Typ(\phi) \wedge Typ(\psi)$.

Now for condition 2. Let $M = \langle \Gamma, \mathcal{V} \rangle$ be a TPL universe, and suppose that $Typ(\phi) \vee Typ(\psi)$ is true in M . Take any $w \in \Gamma$. Since $Typ(\phi \vee \psi)$ is true in M , $Typ(\phi) \vee Typ(\psi)$ is true at w . So either $Typ(\phi)$ is true at w or $Typ(\psi)$ is true at w . Without loss of generality, suppose $Typ(\phi)$ is true at w . Then by the recursive definition of truth in Section 2.2, there is a set $X \in \mathcal{S}$ such that $X \subseteq \Gamma_{\phi}$ and $\frac{\tau(\Gamma \setminus X)}{\tau(\Gamma)} < \epsilon$. For each $w' \in \Gamma_{\phi}$, $w' \in \Gamma_{\phi \vee \psi}$, since if ϕ is true at w' then $\phi \vee \psi$ is true at w' . Therefore, $\Gamma_{\phi} \subseteq \Gamma_{\phi \vee \psi}$. It follows that $X \subseteq \Gamma_{\phi \vee \psi}$. Therefore, by the recursive definition of truth in Section 2.2, $Typ(\phi \vee \psi)$ is true at w . So Lemma 1 implies that $Typ(\phi \vee \psi)$ is true at each world in M . Therefore, $Typ(\phi \vee \psi)$ is true in M . The same conclusion results if $Typ(\psi)$, rather than $Typ(\phi)$, is true at w . And therefore, $Typ(\phi) \vee Typ(\psi) \models Typ(\phi \vee \psi)$.

□

The reverse implications do not always hold: that is, $Typ(\phi) \wedge Typ(\psi) \not\models Typ(\phi \wedge \psi)$, and $Typ(\phi \vee \psi) \not\models Typ(\phi) \vee Typ(\psi)$. To see why $Typ(\phi) \wedge Typ(\psi)$ does not imply $Typ(\phi \wedge \psi)$, consider a TPL universe $M = \langle \Gamma, \mathcal{V} \rangle$ that satisfies the following four conditions. First, Γ contains 100 worlds. Second, $\epsilon = \frac{1}{10}$ and τ is the counting measure. Third, suppose that sentence letter p is true at 91 worlds. Fourth, suppose that sentence letter q is true at exactly 90 of the worlds at which p is true, and q is true at exactly one of the worlds at which p is

false. Then $Typ(p)$ is true at each world in Γ : this follows from the fact that Γ_p contains exactly 91 worlds, and so $\frac{\tau(\Gamma \setminus \Gamma_p)}{\tau(\Gamma)} = \frac{9}{100} < \frac{1}{10} = \epsilon$. For analogous reasons, $Typ(q)$ is true at each world in Γ . Therefore, $Typ(p) \wedge Typ(q)$ is true in M . But $\Gamma_{p \wedge q}$ contains exactly 90 worlds, so $\frac{\tau(\Gamma \setminus \Gamma_{p \wedge q})}{\tau(\Gamma)} = \frac{10}{100} = \frac{1}{10} \not\leq \epsilon$. It follows that $Typ(p \wedge q)$ is not true at any world in M , and thus, that $Typ(p \wedge q)$ is false in M . Therefore, $Typ(\phi) \wedge Typ(\psi) \not\equiv Typ(\phi \wedge \psi)$.

To see why $Typ(\phi \vee \psi)$ does not imply $Typ(\phi) \vee Typ(\psi)$, consider a TPL universe $M = \langle \Gamma, \mathcal{V} \rangle$ that satisfies the following four conditions. First, Γ contains 100 worlds. Second, $\epsilon = \frac{1}{10}$ and τ is the counting measure. Third, suppose that sentence letter p is true at 50 worlds. Fourth, suppose that sentence letter q is true at all and only the worlds at which p is false. Then $Typ(p \vee q)$ is true at each world in Γ : this follows from the fact that $\Gamma_{p \vee q}$ contains all 100 worlds, and so $\frac{\tau(\Gamma \setminus \Gamma_{p \vee q})}{\tau(\Gamma)} = \frac{0}{100} = 0 < \epsilon$. But Γ_p contains exactly 50 worlds, and so $\frac{\tau(\Gamma \setminus \Gamma_p)}{\tau(\Gamma)} = \frac{50}{100} = \frac{1}{2} \not\leq \epsilon$. It follows that $Typ(p)$ is false at each world in Γ . For analogous reasons, $Typ(q)$ is false at each world in Γ . Thus, $Typ(p) \vee Typ(q)$ is false at each world in Γ , so $Typ(p) \vee Typ(q)$ is false in M . And therefore, $Typ(\phi \vee \psi) \not\equiv Typ(\phi) \vee Typ(\psi)$.

2.5.2 Deduction Theorems

The deduction theorem of propositional logic says that for any finite set of well-formed formulas Σ , and for any well-formed formulas ϕ and ψ , $\Sigma \vdash \phi \rightarrow \psi$ if and only if $\Sigma, \phi \vdash \psi$. The semantic deduction theorem of propositional logic is similar: for any finite set of well-formed formulas Σ , and for any well-formed formulas ϕ and ψ , $\Sigma \models \phi \rightarrow \psi$ if and only if $\Sigma, \phi \models \psi$.

The left-to-right direction of each biconditional holds in TPL. The right-to-left direction of each biconditional, however, does not. Let us see why.

To start, consider the deduction theorem for \vdash .

Theorem 4 (Partial Deduction Theorem). *Let Σ be a finite set of well-formed formulas of TPL. Let ϕ and ψ be well-formed formulas of TPL. If $\Sigma \vdash \phi \rightarrow \psi$ then $\Sigma, \phi \vdash \psi$.*

Proof. Let $\Sigma = \{\sigma_1, \sigma_2, \dots, \sigma_n\}$. Suppose $\Sigma \vdash \phi \rightarrow \psi$. Then there is a closed tree whose

initial node consists of the lines σ_i, a ($1 \leq i \leq n$) along with the line $\neg(\phi \rightarrow \psi), w$; call this tree \mathbf{p} .

Consider a tree \mathbf{p}' whose initial node consists of the lines σ_i, a ($1 \leq i \leq n$) along with the line ϕ, a and the line $\neg\psi, w$. To complete the proof of this theorem, we explain how to decompose these initial lines of \mathbf{p}' to close each branch. The decompositions depend on exactly how branches are closed in \mathbf{p} . Put roughly, any way that \mathbf{p} reaches a contradiction can be reached by \mathbf{p}' , because the decomposition of the line $\neg(\phi \rightarrow \psi), w$ is already ‘built into’ the initial node of \mathbf{p}' . But to be completely rigorous, we must check all the possible cases: we must check all the ways that a branch of \mathbf{p} might use the line $\neg(\phi \rightarrow \psi), w$ to reach a contradiction.

Suppose \mathbf{p} reaches a contradiction on a branch b without decomposing the line $\neg(\phi \rightarrow \psi), w$. Then either \mathbf{p} reaches that contradiction by producing a line of the form $\phi \rightarrow \psi, w$ from the σ_i, a lines, or \mathbf{p} reaches that contradiction without using the line $\neg(\phi \rightarrow \psi), w$ at all. In the latter case, the same contradiction can be reached on a corresponding branch b' in \mathbf{p}' by applying the exact same rules in the exact same way. In the former case, it is possible to create a pair of branches in \mathbf{p}' that correspond to b , both of which close, by doing the following: apply the exact same rules in the exact same way as in \mathbf{p} to derive $\phi \rightarrow \psi, w$; use the conditional rule to create a $\neg\phi, w$ branch b_1 and a ψ, w branch b_2 ; use the second all-to-less index rule to decompose the line ϕ, a in the initial node to ϕ, w on b_1 . Note that both of the newly-created branches contain a contradiction: b_1 contains lines ϕ, w and $\neg\phi, w$, and b_2 contains lines ψ, w and $\neg\psi, w$.²⁰ So both branches close.

Alternatively, suppose \mathbf{p}' reaches a contradiction on a branch b by decomposing the line $\neg(\phi \rightarrow \psi), w$ into the pair of lines ϕ, w and $\neg\psi, w$ via the ‘world’ index negated conditional rule. Then the same contradiction can be reached on a corresponding branch b' in \mathbf{p}' , in the following way. First, apply the exact same rules in the exact same way as in \mathbf{p} – apart from that decomposition of initial line $\neg(\phi \rightarrow \psi), w$ of course, since that initial line in \mathbf{p} is not

²⁰The line $\neg\psi, w$ is in the initial node of \mathbf{p}' .

an initial line in \mathfrak{p}' . Second, apply the second all-to-less index rule to decompose the line ϕ, a in the initial node of \mathfrak{p}' to get the line ϕ, w . Then b' has both ϕ, w and $\neg\psi, w$ on it, just as b does. So the contradiction is reached on b' in just the same way that it is reached on b .

□

The other direction of this theorem does not hold: it is not the case that if $\Sigma, \phi \vdash \psi$ then $\Sigma \vdash \phi \rightarrow \psi$. A counterexample has already been given. As shown in Section 2.3, $p \vdash Typ(p)$, but $\not\vdash p \rightarrow Typ(p)$.

Let us now consider the deduction theorem for \models .

Theorem 5 (Partial Semantic Deduction Theorem). *Let σ be a finite set of well-formed formulas of TPL. Let ϕ and ψ be well-formed formulas of TPL. If $\Sigma \models \phi \rightarrow \psi$ then $\Sigma, \phi \models \psi$.*

Proof. Let $\Sigma = \{\sigma_1, \sigma_2, \dots, \sigma_n\}$. Let $M = \langle \Gamma, \mathcal{V} \rangle$ be a TPL universe such that $\sigma_1, \sigma_2, \dots, \sigma_n$, and ϕ are true in M . Suppose $\Sigma \models \phi \rightarrow \psi$. Then $\phi \rightarrow \psi$ is true at each world in Γ . Take $w \in \Gamma$. Then ϕ and $\phi \rightarrow \psi$ are both true at w . Therefore, ψ is true at w as well. Since this holds for each $w \in \Gamma$, ψ is true in M . And since this holds for arbitrary M , it follows that $\Sigma, \phi \models \psi$.

□

The other direction of this theorem does not hold: it is not the case that if $\Sigma, \phi \models \psi$ then $\Sigma \models \phi \rightarrow \psi$. For example, note that $p \models Typ(p)$: this holds because for any TPL universe $M = \langle \Gamma, \mathcal{V} \rangle$, if $\models_M p$ —that is, if p is true at each $w \in \Gamma$ —then $\{w \in \Gamma \mid \models_{M,w} p\} = \Gamma$, and so $Typ(p)$ is true at each $w \in \Gamma$. But $\not\models p \rightarrow Typ(p)$. To see why, let Γ be a large set, let $M = \langle \Gamma, \mathcal{V} \rangle$ be an TPL universe, and suppose p is true at just one world $w \in \Gamma$. Then it is not the case that $Typ(p)$ is true in M . So $\models_{M,w} p$ but $\not\models_{M,w} Typ(p)$. Therefore, $\not\models_{M,w} p \rightarrow Typ(p)$, and so $\not\models p \rightarrow Typ(p)$.

2.5.3 Modal Logic and TPL

In this subsection, we briefly discuss an important difference between TPL and modal logic. The typicality operator Typ is substantially different from the necessity operator \Box . Modal logics generally adopt the K -axiom schema: $\Box(A \rightarrow B) \rightarrow (\Box A \rightarrow \Box B)$. The corresponding schema for TPL would be: $Typ(A \rightarrow B) \rightarrow (Typ(A) \rightarrow Typ(B))$. But it can be shown that instances of this schema are false in some TPL universes. Thus, TPL is different in kind from the main systems of modal logic.

As an example of the falsity of schema $Typ(A \rightarrow B) \rightarrow (Typ(A) \rightarrow Typ(B))$ in some cases, consider the well-formed formula $Typ(p \rightarrow q) \rightarrow (Typ(p) \rightarrow Typ(q))$. And consider the following TPL universe $\langle \Gamma, \mathcal{V} \rangle$. Γ contains 100 worlds. At 91 of those worlds, p is true. Of those 91 worlds, q is true at exactly 82 of them. In addition, q is false at each of the 9 worlds at which p is false. From all this, it follows that $p \rightarrow q$ is true at exactly 91 worlds: it is true at each of the 82 worlds at which q is true, it is true at each of the 9 worlds at which p is false, and it is false at the 9 worlds where p is true but q is false. Let $\epsilon = \frac{1}{10}$, and let τ be the counting measure. Then $\frac{\tau(\Gamma \setminus \Gamma_{p \rightarrow q})}{\tau(\Gamma)} = \frac{9}{100} < \frac{1}{10} = \epsilon$, so $Typ(p \rightarrow q)$ is true at each world in Γ . Similarly, $\frac{\tau(\Gamma \setminus \Gamma_p)}{\tau(\Gamma)} = \frac{9}{100} < \frac{1}{10} = \epsilon$, so $Typ(p)$ is true at each world in Γ . But $\frac{\tau(\Gamma \setminus \Gamma_q)}{\tau(\Gamma)} = \frac{18}{100} \not< \frac{1}{10} = \epsilon$, so $Typ(q)$ is false at each world in Γ . Therefore, $Typ(p \rightarrow q) \rightarrow (Typ(p) \rightarrow Typ(q))$ is false in $\langle \Gamma, \mathcal{V} \rangle$.

This is a feature of TPL, not a bug. The K -axiom makes intuitive sense in modal logic: intuitively, it says that if $p \rightarrow q$ holds at each world, and if p holds at each world, then q holds at each world. But the corresponding axiom for typicality is overly strong, since typical statements need not hold everywhere. Indeed, typical statements are generally liable to exceptions: typicality is *nearly* all, not *absolutely* all. So if $p \rightarrow q$ is typical and p is typical, it follows that $p \rightarrow q$ is true at nearly all worlds, and it follows that p is true at nearly all worlds. But since $p \rightarrow q$ and p may not be true at all the same worlds, it does not follow that q is true at nearly all worlds. So there may not be enough q worlds for q to be typical.

3 Typicality Intuitionistic Logic

In this section, we describe an intuitionistic logic for typicality, which we call ‘typicality intuitionistic logic’, or TIL. To contrast TIL with TPL, we note at the outset a shift in perspective from the impersonal and objective interpretation of TPL using sets of possible worlds (Section 2) to the personal and subjective interpretation of TIL using sets of ‘credal states’ (below).²¹ It is important not to confuse the credal states discussed in this section with the physical states (i.e., quantum states, microstates, or generically ‘states of the world’) discussed in the previous. Whereas the states of the previous section are objective properties, those of the current reflect subjective dispositions toward claims. These dispositions guide an agent’s judgment about what is true and what is typical.

Recall that in TPL, we associate each well-formed formula ϕ to the set of possible worlds Γ_ϕ at which ϕ holds. In this interpretation, when a statement holds at a world, it is a *fact* of that world, and the formalism of TPL makes precise how the notion of typicality can be integrated into such an outlook. In particular, if a statement holds at a sufficiently large subset of possible worlds, then the typicality of that statement is a *fact* of every world in the universe.

In TIL, on the other hand, when a credence assigns maximal belief to a statement, an agent with that credence is justified in making a *judgment* that the claim holds, regardless of whether or not it actually does. Thus, in the alternative formalism given below, when a credence assigns a sufficiently large degree of belief to a statement, an agent with that credence is justified in make a *judgment* that the claim is typical. The validity of all such judgments is determined not by the actual states of the world, which are unknown to the subject making the judgments, but rather by the *context* within which the judgment is being made. In proposed formalism, we interpret this notion of context as a constraint on the

²¹For our purposes, a credal state is a probability space $(\Omega, \mathcal{F}, \nu)$ with ν representing the credences of a rational agent regarding the measurable subsets in \mathcal{F} . It is assumed for the purposes of this discussion that each subset in \mathcal{F} represents a proposition in TPL.

possible credal states of the agent making the judgment. The semantic difference between possible worlds and possible credal states in the two formalisms mirrors the philosophical distinction between ‘typicality facts’ and ‘typicality judgments’.

3.1 The Language

To define the language of TIL we extend a fragment of the non-dependent version of intuitionistic Martin-Löf type theory (MLTT) (Martin-Löf, 1985) by specifying rules for an additional typicality type former **Typ**. The syntactic rules for the **Typ**-type specified below are a modification of an earlier development by Crane to append a probability type (denoted **Prob**) to the syntax of MLTT. See Crane (2018) for further details.

For clarity, we reserve capital Greek, lowercase Greek, and lowercase Roman letters to represent the different primitive notions of context, types, and terms, respectively. The basic components of the language are called judgments, each of which has the form of one of the following three primitive expressions:²²

Formal	Natural language description
$\Delta \text{ ctx}$	Δ is a well-formed context
$\phi : \mathbf{Type}$	ϕ is a type
$a : \phi$	a is a term of type ϕ

Importantly, the syntax of MLTT does not permit ‘untyped’ statements, e.g., it is meaningless to refer to a term without reference to its type. In the preferred propositions-as-types interpretation we appeal to below, a judgment of the form $\phi : \mathbf{Type}$ is interpreted to mean that ϕ is a proposition, and $a : \phi$ is interpreted to mean that a is a proof of the proposition ϕ (Curry and Feys, 1959; Howard, 1969). When understood in this light, the rules of MLTT presented below can be understood as an algorithmic prescription for how to prove compound

²²The full syntax of MLTT has two additional primitive notions of ‘judgmental equality’ which play no substantive role in our treatment below and are thus omitted.

propositions, as we illustrate below in the case of the product type.²³

In MLTT, and thus in TIL, all judgments are made in a specific context and are expressed in the form

$$\begin{array}{c} \text{context} \vdash \text{judgment} \\ \Delta \vdash \mathcal{J}, \end{array} \quad (1)$$

where \mathcal{J} has the form of one of the above primitive judgments and Δ is a finite list of the form

$$a_1 : \phi_1, \dots, a_n : \phi_n.$$

The statement in (1) can be interpreted pre-formally to mean that the judgment \mathcal{J} on the right is made in the context Δ on the left, where a context consists of a string of prior judgments about types ϕ_1, \dots, ϕ_n .

From these primitive elements, the syntax of MLTT is built by specifying a collection of rules for how to derive new judgments from old. In the non-dependent version of MLTT featured here, these derived judgments correspond to the non-logical vocabulary for constructing the \times -type, $+$ -type, and $\mathbf{0}$ -type. For example, the product (\times) type is defined by the following rules.

$$\frac{\Delta \vdash \phi : \mathbf{Type} \quad \Delta \vdash \psi : \mathbf{Type}}{\Delta \vdash \phi \times \psi : \mathbf{Type}} \quad (\times\text{-form})$$

This formation rule says that given two types ϕ and ψ a new type can be formed, denoted $\phi \times \psi$. This is akin to saying that $\phi \times \psi$ is a ‘well-formed formula’ in the language of MLTT, or more appropriately that $\phi \times \psi$ is a ‘well-formed type’ in the context Δ . In the propositions-as-types interpretation, this rule corresponds to the conjunction formation rule in proposition logic, by which two well-formed formulas ϕ and ψ in PL can be combined to another well-formed formula $\phi \wedge \psi$ in PL.

²³With this understanding of ‘terms’ as ‘proofs’, the meaninglessness of untyped statements, e.g., let ‘ a ’ be a term, becomes apparent; for if we interpret terms as proofs, the statement “let a be a proof” leaves ambiguous what it is that a is a proof of.

$$\begin{array}{c}
\Delta \vdash \phi : \mathbf{Type} \quad \Delta \vdash \psi : \mathbf{Type} \\
\hline
\Delta \vdash a : \phi \quad \Delta \vdash b : \psi \\
\hline
\Delta \vdash (a, b) : \phi \times \psi \quad (\times\text{-intro})
\end{array}$$

This introduction rule describes how terms of the product type $\phi \times \psi$ are constructed (or introduced) by combining terms $a : \phi$ and $b : \psi$ to form $(a, b) : \phi \times \psi$. In the corresponding propositions-as-types interpretation this rule establishes the truth conditions for $\phi \wedge \psi$. In particular, to verify that $\phi \wedge \psi$ is true, verify that ϕ is true and ψ is true individually. Combining these two individual verifications serves as a verification of their conjunction.

$$\begin{array}{c}
\Delta, a : \phi, b : \psi \vdash \rho : \mathbf{Type} \\
\hline
\Delta, a : \phi, b : \psi \vdash d(a, b) : \rho \\
\hline
\Delta, z : \phi \times \psi \vdash \mathbf{split}_d(z) : \rho \quad (\times\text{-elim})
\end{array}$$

This elimination rule asserts that to define a function out of $\phi \times \psi$ it is sufficient to define how the function acts on pairs of the form (a, b) for $a : \phi$ and $b : \psi$. Intuitively, though somewhat informally, this rule states implicitly that all terms of $\phi \times \psi$ consist of pairs of the form (a, b) for $a : \phi$ and $b : \psi$.²⁴ ²⁵ In the corresponding propositions-as-types interpretation, the elimination rule describes the conditions under which it is justified to deduce ρ on the basis of $\phi \wedge \psi$. In particular, if ρ holds whenever ϕ and ψ both hold, then ρ holds whenever

²⁴To understand the notation $d(a, b) : \rho$ it may be helpful to think of d as a function that takes terms $a : \phi$ and $b : \psi$ as input and returns a term $d(a, b) : \rho$ as output. But in general this is simply notation that denotes a new term ‘ d ’ that is constructed in a way that depends on $a : \phi$ and $b : \psi$. Similarly, and informally, \mathbf{split}_d in the conclusion can be thought of as a function that takes a term $z : \phi \times \psi$ of the product type as input and outputs $\mathbf{split}_d(z) : \rho$.

²⁵In addition to the above three rules, the standard rules for MLTT specify the following computation rule:

$$\begin{array}{c}
\Delta \vdash \phi : \mathbf{Type} \quad \Delta \vdash \psi : \mathbf{Type} \\
\Delta, a : \phi, b : \psi \vdash \rho : \mathbf{Type} \\
\Delta, a : \phi, b : \psi \vdash d(a, b) : \rho \\
\hline
\Delta, (a, b) : \phi \times \psi \vdash \mathbf{split}_d((a, b)) \equiv d(a, b) : \rho \quad (\times\text{-comp})
\end{array}$$

The computation rule establishes coherence between \mathbf{split}_d and the function d from which it is derived. In particular, it establishes that the introduction and elimination rules commute. Because we specialize here to the propositions-as-types case, in which all inhabited types have exactly one unique term up to identity, these computation rules play no material role and are therefore omitted.

$\phi \wedge \psi$ holds.

The coproduct (+) and **0** types are defined by analogous rules of formation, introduction, and elimination; see, e.g., Appendix A of Lumsdaine and Kupulkin (2014) or Tsementzis (2018) for a detailed description of those rules. We also list these rules explicitly in our soundness proof for TIL (Theorem 6).

3.2 The Typicality Type

To build a notion of typicality on top of the existing machinery of MLTT, we define the following rules for a new type **Typ** as follows.

The first formation rule (**Typ-form**) says that if ϕ is a well-formed type in context Δ , then **Typ**(ϕ) is a well-formed type in that context. This operation is akin to the step taken when defining the well-formed formulas of TPL in Section 2, in which $Typ(\phi)$ is a well-formed formula whenever ϕ is a well-formed formula. Expressed formally, the formation rule reads:

$$\frac{\Delta \vdash \phi : \mathbf{Type}}{\Delta \vdash \mathbf{Typ}(\phi) : \mathbf{Type}} \quad (\mathbf{Typ-form})$$

The next introduction rule (**Typ-intro**) says that a judgment that a proposition ϕ is typical can be constructed from any proof of ϕ . In particular, from a proof $a : \phi$ that ϕ holds, we construct $\tau_\phi(a) : \mathbf{Typ}(\phi)$, where τ_ϕ is the constructor for the **Typ**(ϕ)-type.

$$\frac{\Delta \vdash \phi : \mathbf{Type}}{\Delta, a : \phi \vdash \tau_\phi(a) : \mathbf{Typ}(\phi)} \quad (\mathbf{Typ-intro})$$

The elimination rule (**Typ-elim**) says that if ϕ implies ψ (i.e., if every proof of ϕ ($a : \phi$) can be turned into a proof of ψ , i.e., $d(a) : \psi$), then **Typ**(ϕ) implies **Typ**(ψ) (i.e., any justification that ϕ is typical, i.e., $x : \mathbf{Typ}(\phi)$, can be used to construct a justification that

ψ is typical ($\mathbf{imp}_d(x) : \mathbf{Typ}(\psi)$)).²⁶

$$\frac{\Delta \vdash \phi : \mathbf{Type} \quad \Delta \vdash \psi : \mathbf{Type} \quad \Delta, a : \phi \vdash d(a) : \psi}{\Delta, x : \mathbf{Typ}(\phi) \vdash \mathbf{imp}_d(x) : \mathbf{Typ}(\psi)} \quad (\mathbf{Typ}\text{-elim})$$

Finally, the **Typ-0** rule states that a justification of the typicality of the empty type $\mathbf{0}$, i.e., $x : \mathbf{Typ}(\mathbf{0})$, can be used to construct an element of the empty type, i.e., $\sigma(x) : \mathbf{0}$. In relating this rule to TPL, we interpret $\mathbf{0}$ as \perp , so that this rule can be translated to mean that $\mathit{Typ}(\perp)$ implies \perp in any context.

$$\frac{\Delta \text{ ctx}}{\Delta, x : \mathbf{Typ}(\mathbf{0}) \vdash \sigma(x) : \mathbf{0}} \quad (\mathbf{Typ}\text{-0})$$

3.3 The Semantics of TIL

We define a semantics for the above system of rules by associating judgments to subsets of probability spaces, where each probability space is understood as the credal state of a rational agent. First, let Π be the set of all sentences in PL and let Π^* consist of all sentences in TPL, obtained recursively by adding the unary predicate Typ to PL, as in Section 2.²⁷

To formalize this, we let Ω be a set (i.e., the set of possible worlds) and associate each proposition $\phi \in \Pi^*$ to a subset of worlds at which ϕ holds, denoted $\tilde{\phi} \subseteq \Omega$. From this, we have the following:

1. If ϕ and ψ are well-formed formulas, then

$$\widetilde{\phi \wedge \psi} = \{\omega \in \Omega \mid \omega \in \tilde{\phi} \text{ and } \omega \in \tilde{\psi}\} \equiv \tilde{\phi} \cap \tilde{\psi}.$$

²⁶In the **Typ**-elimination rule, \mathbf{imp}_d is a constructor for $\mathbf{Typ}(\psi)$ built from the constructor $d(-)$ in the premises. The main content of the elimination rule is its assertion of the validity of constructing $\mathbf{imp}_d : \mathbf{Typ}(\phi) \rightarrow \mathbf{Typ}(\psi)$ from $d : \phi \rightarrow \psi$.

²⁷In Section 2, we defined the syntax of TPL using only the logical connectives \neg and \rightarrow . We can define the connectives \wedge and \vee in the usual way and supplement TPL with these connectives.

2. If ϕ and ψ are well-formed formulas, then

$$\widetilde{\phi \vee \psi} = \{\omega \in \Omega \mid \omega \in \tilde{\phi} \text{ or } \omega \in \tilde{\psi}\} \equiv \tilde{\phi} \cup \tilde{\psi}.$$

3. If ϕ is a well-formed formula, then

$$\widetilde{\neg\phi} = \{\omega \in \Omega \mid \omega \notin \tilde{\phi}\} \equiv \tilde{\phi}^c.$$

4. If ϕ is a well-formed formula, then we require $\widetilde{Typ(\phi)} \supseteq \tilde{\phi}$.

5. Finally we have $\tilde{\perp} = \emptyset$.

Altogether, a *credal state* is a probability space $(\Omega, \mathcal{F}, \nu)$ that assigns credence $\nu(\tilde{\phi})$ to every proposition $\phi \in \Pi^*$ whose representation satisfies $\tilde{\phi} \in \mathcal{F}$. Any ϕ in TPL for which $\phi \notin \mathcal{F}$ is a proposition for which the agent has no credence, and thus suspends judgment. In total, the credal state $(\Omega, \mathcal{F}, \nu)$ specifies the σ -algebra \mathcal{F} of propositions about which an agent in this state has a credence along with a probability measure ν on (Ω, \mathcal{F}) that specifies the agent's credences.

In what follows, we fix a set of possible worlds Ω along with a representation $\tilde{\phi} \subseteq \Omega$ of each $\phi \in \Pi^*$. For any fixed $\epsilon \in (1/2, 1)$, we write

$$\mathcal{P}(\Pi^*) = \{(\Omega, \mathcal{F}, \nu) \mid \forall \phi \in \Pi^* \tilde{\phi} \in \mathcal{F} \Rightarrow \widetilde{Typ(\phi)} \in \mathcal{F}, \forall \phi \in \Pi^* \nu(\widetilde{Typ(\phi)}) = 1 \Leftrightarrow \nu(\tilde{\phi}) \geq 1 - \epsilon\}$$

to denote the set of admissible credal states for a rational agent holding beliefs about propositions in Π^* . For every $\phi \in \Pi^*$ we define

$$\mathcal{S}_\phi = \{(\Omega, \mathcal{F}, \nu) \mid \tilde{\phi} \in \mathcal{F}\}$$

$$\mu_\phi = \{(\Omega, \mathcal{F}, \nu) \in \mathcal{P}(\Pi^*) \mid \tilde{\phi} \in \mathcal{F} \text{ and } \nu(\tilde{\phi}) = 1\}$$

We interpret the syntax from Section 3.1 and 3.2 into the semantics of TIL by regarding all type-theoretic judgments as set-theoretic statements. In particular, we have the following symbolic correspondence for symbols in type theory, propositional logic, and set theory:

MLTT	PL	set theory
$\Delta \text{ ctx}$	—	$\Delta \subseteq \mathcal{P}(\Pi^*)$
$\phi : \mathbf{Type}$	ϕ well-formed formula	\mathcal{S}_ϕ
$a : \phi$	ϕ	μ_ϕ
\vdash	\vdash	\subseteq
$\phi \times \psi$	$\phi \wedge \psi$	$\tilde{\phi} \cap \tilde{\psi}$
$\phi + \psi$	$\phi \vee \psi$	$\tilde{\phi} \cup \tilde{\psi}$
$\mathbf{0}$	\perp	\emptyset

In type theory, on the left side of the turnstile we interpret commas as \cap . With this translation, the basic judgments of MLTT are interpreted as:

Syntax	Semantics
$\Delta \text{ ctx}$	$\Delta \subseteq \mathcal{P}(\Pi^*)$
$\Delta \vdash \phi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\phi$
$\Delta \vdash a : \phi$	$\Delta \subseteq \mu_\phi$

Thus, the syntax $\Delta \text{ ctx}$ translates to $\Delta \subseteq \mathcal{P}(\Pi^*)$, thus justifying our interpretation of the context as a constraint on the credal states in which a judgment is made. The initial (‘empty’) context \bullet is thus the one without any constraints on the credal state, namely $\bullet \equiv \mathcal{P}(\Pi^*)$. The syntax $\Delta \vdash \phi : \mathbf{Type}$ translates to $\Delta \subseteq \mathcal{S}_\phi$, which imposes the constraint that the admissible credal states are those which assign some credence to ϕ . The syntax $\Delta \vdash a : \phi$ translates to $\Delta \subseteq \mu_\phi$, which imposes the constraint that the admissible credal states are those which assign maximal credence to ϕ .

For example, an agent in a specific credal state has a σ -algebra \mathcal{F} corresponding to a subset of the propositions in Π^* and a probability measure ν that assigns a credence to each measurable subset of \mathcal{F} . If the subset $\tilde{\phi}$ corresponding to proposition ϕ is a measurable subset of \mathcal{F} , then the agent having this credence would make the judgment that ‘ ϕ holds’ only if $\nu(\tilde{\phi}) = 1$ and that ‘ ϕ is typical’ only if $\nu(\tilde{\phi}) \geq 1 - \epsilon$. The semantic interpretation of the rules of TIL does not require an agent’s credal state to be pinned down to a single

probability space in order for a judgment to be justified. The semantics only require that the agent’s credal state lies in a subset of possible credal states that are consistent with the given judgment.

To illustrate the semantic translation of the syntax, a deduction of the form

$$\begin{array}{ccc}
 \Delta \text{ ctx} & & \Delta \subseteq \mathcal{P}(\Pi^*) \\
 \Delta \vdash \phi : \mathbf{Type} & & \Delta \subseteq \mathcal{S}_\phi \\
 \frac{\Delta, a : \phi \vdash \psi : \mathbf{Type}}{\Delta, a : \phi \vdash b : \psi} & \text{translates to} & \frac{\Delta \cap \mu_\phi \subseteq \mathcal{S}_\psi}{\Delta \cap \mu_\phi \subseteq \mu_\psi}.
 \end{array}$$

3.4 Soundness

Theorem 6. *The syntax of TIL is sound with respect to the above interpretation.*

Proof. To prove soundness, we interpret each of the rules for the \times , $+$, $\mathbf{0}$, and \mathbf{Type} types into the semantics and show that the rule holds. We begin by specifying the interpretation of the rules for contexts.

- **Structural rules, \bullet -ctx:**

Syntax	Semantics
$\frac{}{\bullet \text{ ctx}}$	$\frac{}{\mathcal{P}(\Pi^*) \subseteq \mathcal{P}(\Pi^*)}$

Holds trivially: Every set is a subset of itself.²⁸

- **Structural rules, ext-ctx**

²⁸This rule states that there is an initial ‘empty’ context \bullet . In the semantics, the context places constraints on an agent’s credal states, and thus this initial ‘empty’ context corresponds to a context without constraints, i.e., $\Delta \equiv \mathcal{P}(\Pi^*)$.

Syntax	Semantics
$\frac{\Delta \text{ ctx} \quad \Delta \vdash \phi : \mathbf{Type}}{\Delta, x : \phi \text{ ctx}}$	$\frac{\Delta \subseteq \mathcal{P}(\Pi^*) \quad \Delta \subseteq \mathcal{S}_\phi}{\Delta \cap \mu_\phi \subseteq \mathcal{P}(\Pi^*)}$

By assumption $\Delta \subseteq \mathcal{P}(\Pi^*)$, and thus $\Delta \cap A \subseteq \Delta \subseteq \mathcal{P}(\Pi^*)$ for all other sets A . Instantiating $A = \mu_\phi$ gives the result.

• **Structural rules, ax-ctx**

Syntax	Semantics
$\frac{\Delta, a : \phi, \Xi \text{ ctx}}{\Delta, a : \phi, \Xi \vdash a : \phi}$	$\frac{\Delta \cap \mu_\phi \cap \Xi \subseteq \mathcal{P}(\Pi^*)}{\Delta \cap \mu_\phi \cap \Xi \subseteq \mu_\phi}$

By assumption, $\Delta \cap \mu_\phi \cap \Xi$ is a set, and for any sets A and B it is always the case that $A \cap B \subseteq A$, yielding the result.

It follows from these structural rules for contexts that every context is a finite list of judgments of the form

$$(a_1 : \phi_1, \dots, a_n : \phi_n) \text{ ctx},$$

which in our semantic interpretation translates to

$$\mu_{\phi_1} \cap \dots \cap \mu_{\phi_n} \subseteq \mathcal{P}(\Pi^*).$$

Thus, in our semantic treatment, every context Δ can be expressed in the form

$$\Delta \equiv \mu_{\phi_1} \cap \dots \cap \mu_{\phi_n} \tag{2}$$

for some finite list $\phi_1, \dots, \phi_n \in \Pi^*$. This specific representation will become useful when we prove soundness for the coproduct and typicality types below.

We next prove soundness for the product type.

- **Product type**, formation rule:

Syntax	Semantics
$\Delta \vdash \phi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\phi$
$\Delta \vdash \psi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\psi$
<hr style="width: 100%; border: 0.5px solid black;"/> $\Delta \vdash \phi \times \psi : \mathbf{Type}$	<hr style="width: 100%; border: 0.5px solid black;"/> $\Delta \subseteq \mathcal{S}_{\phi \wedge \psi}$

Let $(\Omega, \mathcal{F}, \nu) \in \Delta$ so that $\tilde{\phi} \in \mathcal{F}$ and $\tilde{\psi} \in \mathcal{F}$. Then $\widetilde{\phi \wedge \psi} \equiv \tilde{\phi} \cap \tilde{\psi} \in \mathcal{F}$ because \mathcal{F} is a σ -algebra on Ω and is closed under intersection. It follows that $\Delta \subseteq \mathcal{S}_{\phi \wedge \psi}$, as claimed.

- **Product type**, introduction rule:

Syntax		Semantics	
$\Delta \vdash \phi : \mathbf{Type}$	$\Delta \vdash a : \phi$	$\Delta \subseteq \mathcal{S}_\phi$	$\Delta \subseteq \mu_\phi$
$\Delta \vdash \psi : \mathbf{Type}$	$\Delta \vdash b : \psi$	$\Delta \subseteq \mathcal{S}_\psi$	$\Delta \subseteq \mu_\psi$
<hr style="width: 100%; border: 0.5px solid black;"/> $\Delta \vdash (a, b) : \phi \times \psi$		<hr style="width: 100%; border: 0.5px solid black;"/> $\Delta \subseteq \mu_{\phi \wedge \psi}$	

First note that $\mu_\phi \subseteq \mathcal{S}_\phi$ and $\mu_\psi \subseteq \mathcal{S}_\psi$, so that the premises together imply $\Delta \subseteq \mu_\phi \cap \mu_\psi$. Now take any $(\Omega, \mathcal{F}, \nu) \in \mu_\phi \cap \mu_\psi$. Since $\tilde{\phi} \in \mathcal{F}$ and $\tilde{\psi} \in \mathcal{F}$ by assumption, we must have $\tilde{\phi} \cap \tilde{\psi} \in \mathcal{F}$ because \mathcal{F} is a σ -algebra. Furthermore, by the equivalence $\neg(\phi \wedge \psi) \equiv \neg\phi \vee \neg\psi$, we have $(\tilde{\phi} \cap \tilde{\psi})^c \equiv \tilde{\phi}^c \cup \tilde{\psi}^c$. The assumption that $\nu(\tilde{\phi}) = \nu(\tilde{\psi}) = 1$ implies $\nu(\tilde{\phi}^c) = \nu(\tilde{\psi}^c) = 0$, and thus

$$\nu(\tilde{\phi} \cap \tilde{\psi}) = 1 - \nu(\tilde{\phi}^c \cup \tilde{\psi}^c) \geq 1 - \nu(\tilde{\phi}^c) - \nu(\tilde{\psi}^c) = 1.$$

It follows that $(\Omega, \mathcal{F}, \nu) \in \mu_{\phi \wedge \psi}$, so that $\Delta \subseteq \mu_{\phi \wedge \psi}$, as claimed.

- **Product type**, elimination rule:

Syntax	Semantics
$\Delta \vdash \phi : \mathbf{Type} \quad \Delta \vdash \psi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\phi \quad \Delta \subseteq \mathcal{S}_\psi$
$\Delta, a : \phi, b : \psi \vdash \rho : \mathbf{Type}$	$\Delta \cap \mu_\phi \cap \mu_\psi \subseteq \mathcal{S}_\rho$
$\Delta, a : \phi, b : \psi \vdash d(a, b) : \rho$	$\Delta \cap \mu_\phi \cap \mu_\psi \subseteq \mu_\rho$
$\Delta, z : \phi \times \psi \vdash \mathbf{split}_d(z) : \rho$	$\Delta \cap \mu_{\phi \wedge \psi} \subseteq \mu_\rho$

Let $(\Omega, \mathcal{F}, \nu) \in \Delta \cap \mu_{\phi \wedge \psi}$. Then, in particular, we must have $(\Omega, \mathcal{F}, \nu) \in \mu_{\phi \wedge \psi}$, from which it follows that $\min(\nu(\tilde{\phi}), \nu(\tilde{\psi})) \geq \nu(\tilde{\phi} \cap \tilde{\psi}) = 1$; whence $\nu(\tilde{\phi}) = \nu(\tilde{\psi}) = 1$ and $(\Omega, \mathcal{F}, \nu) \in \Delta \cap \mu_\phi \cap \mu_\psi$. Now by assumption, we have $\Delta \cap \mu_\phi \cap \mu_\psi \subseteq \mu_\rho$ so that $(\Omega, \mathcal{F}, \nu) \in \mu_\rho$, and thus $\Delta \cap \mu_{\phi \wedge \psi} \subseteq \mu_\rho$, as claimed.

Before we move on to discuss the coproduct type, we can use the rules for product type to deduce that $\mu_\phi \cap \mu_\psi = \mu_{\phi \wedge \psi}$ for any $\phi, \psi \in \Pi^*$. From this and the representation of contexts in the form (2), we can equivalently express any context as

$$\Delta \equiv \mu_{\phi_1 \wedge \dots \wedge \phi_n}, \tag{3}$$

which can more compactly be written as

$$\Delta \equiv \mu_\Phi$$

for some $\Phi \in \Pi^*$, because $\tilde{\phi}_1 \cap \dots \cap \tilde{\phi}_n \in \mathcal{F}$ whenever $\tilde{\phi}_1, \dots, \tilde{\phi}_n \in \mathcal{F}$. This representation plays a role in our proof of soundness for the coproduct elimination rule below.

We next discuss the coproduct type.

- **Coproduct type**, formation rule:

Syntax	Semantics
$\Delta \vdash \phi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\phi$
$\Delta \vdash \psi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\psi$
$\Delta \vdash \phi + \psi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_{\phi \vee \psi}$

Let $(\Omega, \mathcal{F}, \nu) \in \mathcal{S}_\phi \cap \mathcal{S}_\psi$, then $\tilde{\phi}^c \in \mathcal{F}$, $\tilde{\psi}^c \in \mathcal{F}$, and $\tilde{\phi} \cap \tilde{\psi} \in \mathcal{F}$, because \mathcal{F} is an algebra. Finally, by definition we have $\widetilde{\phi \vee \psi} \equiv \tilde{\phi} \cup \tilde{\psi} \equiv (\tilde{\phi}^c \cap \tilde{\psi}^c)^c \in \mathcal{F}$, because \mathcal{F} is a σ -algebra and is closed under complementation and countable intersection. It follows that $(\Omega, \mathcal{F}, \nu) \in \mathcal{S}_{\phi \vee \psi}$.

- **Coproduct type**, introduction rule 1:

Syntax	Semantics
$\Delta \vdash \phi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\phi$
$\Delta \vdash \psi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\psi$
$\Delta \vdash a : \phi$	$\Delta \subseteq \mu_\phi$
<hr style="width: 100%; border: 0.5px solid black;"/>	<hr style="width: 100%; border: 0.5px solid black;"/>
$\Delta \vdash \mathbf{inl}(a) : \phi + \psi$	$\Delta \subseteq \mu_{\phi \vee \psi}$

The three assumptions combine to imply $\Delta \subseteq \mu_\phi \cap \mathcal{S}_\psi$, so that any $(\Omega, \mathcal{F}, \nu) \in \Delta$ satisfies $\tilde{\phi}, \tilde{\psi} \in \mathcal{F}$ and $\nu(\tilde{\phi}) = 1$. Because \mathcal{F} is an algebra (or alternatively by the preceding formation rule), we have $\tilde{\phi} \cup \tilde{\psi} \in \mathcal{F}$, and so ν assigns measure to it, and since probability measures are increasing we must have $\nu(\tilde{\phi} \cup \tilde{\psi}) \geq \nu(\tilde{\phi}) = 1$; whence $\nu \in \mu_{\phi \vee \psi}$, as claimed.

- **Coproduct type**, introduction rule 2:

Syntax	Semantics
$\Delta \vdash \phi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\phi$
$\Delta \vdash \psi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\psi$
$\Delta \vdash b : \psi$	$\Delta \subseteq \mu_\psi$
<hr style="width: 100%; border: 0.5px solid black;"/>	<hr style="width: 100%; border: 0.5px solid black;"/>
$\Delta \vdash \mathbf{inr}(b) : \phi + \psi$	$\Delta \subseteq \mu_{\phi \vee \psi}$

The three assumptions combine to imply $\Delta \subseteq \mu_\psi \cap \mathcal{S}_\phi$, so that any $(\Omega, \mathcal{F}, \nu) \in \Delta$ satisfies $\tilde{\phi}, \tilde{\psi} \in \mathcal{F}$ and $\nu(\tilde{\psi}) = 1$. Because \mathcal{F} is an algebra (or alternatively by the preceding formation rule), we have $\tilde{\phi} \cup \tilde{\psi} \in \mathcal{F}$, and so ν assigns measure to it, and since probability measures are increasing we must have $\nu(\tilde{\phi} \cup \tilde{\psi}) \geq \nu(\tilde{\psi}) = 1$; whence $\nu \in \mu_{\phi \vee \psi}$, as claimed.

- **Coproduct type**, elimination rule:

Syntax	Semantics
$\Delta \vdash \phi : \mathbf{Type} \quad \Delta \vdash \psi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\phi \quad \Delta \subseteq \mathcal{S}_\psi$
$\Delta, a : \phi, b : \psi \vdash \rho : \mathbf{Type}$	$\Delta \cap \mu_\phi \cap \mu_\psi \subseteq \mathcal{S}_\rho$
$\Delta, a : \phi \vdash d_l(a) : \rho$	$\Delta \cap \mu_\phi \subseteq \mu_\rho$
$\Delta, b : \psi \vdash d_r(b) : \rho$	$\Delta \cap \mu_\psi \subseteq \mu_\rho$
$\Delta, z : \phi + \psi \vdash \mathbf{case}_{d_l, d_r}(z) : \rho$	$\Delta \cap \mu_{\phi \vee \psi} \subseteq \mu_\rho$

Here we use the representation in (3) to express $\Delta \equiv \mu_\Phi$ for some $\Phi \subseteq \Pi^*$, so that the third and fourth assumptions and the conclusion on the righthand side, respectively, become

$$\mu_{\Phi \wedge \phi} \subseteq \mu_\rho,$$

$$\mu_{\Phi \wedge \psi} \subseteq \mu_\rho \quad \text{and}$$

$$\mu_{\Phi \wedge (\phi \vee \psi)} \subseteq \mu_\rho.$$

By assumption, we have $\mu_{\Phi \wedge \phi} \subseteq \mu_\rho$. Thus, any ν that satisfies $\nu(\tilde{\Phi} \cap \tilde{\phi}) = 1$ must also satisfy $\nu(\tilde{\rho}) = 1$, which is possible only if $\tilde{\Phi} \cap \tilde{\phi} \subseteq \tilde{\rho}$. For suppose that there is some $\omega \in \tilde{\Phi} \cap \tilde{\phi}$ for which $\omega \notin \tilde{\rho}$. Then there is a measurable space $(\Omega, \mathcal{F}_\omega, \nu_\omega)$ with powerset σ -algebra \mathcal{F}_ω on Ω and ν_ω the atomic measure at $\{\omega\}$ (i.e., $\nu_\omega(\{\omega\}) = 1$). With $\omega \in \tilde{\Phi} \cap \tilde{\phi}$, it follows that

$\nu_\omega(\tilde{\Phi} \cap \tilde{\phi}) \geq \nu_\omega(\{\omega\}) = 1$ and $\nu_\omega(\tilde{\rho}) = 0$, contradicting the assumption. By applying an analogous argument to the fourth assumption, we must have $\tilde{\Phi} \cap \tilde{\phi} \subseteq \tilde{\rho}$.

For the conclusion, note that $\tilde{\Phi} \cap (\tilde{\phi} \cup \tilde{\psi}) \equiv (\tilde{\Phi} \cap \tilde{\phi}) \cup (\tilde{\Phi} \cap \tilde{\psi})$, so that the conclusion reads

$$\mu_{(\Phi \wedge \phi) \vee (\Phi \wedge \psi)} \subseteq \mu_\rho.$$

Now, suppose $(\Omega, \mathcal{F}, \nu) \in \mu_{(\Phi \wedge \phi) \vee (\Phi \wedge \psi)}$ so that $\nu((\tilde{\Phi} \cap \tilde{\phi}) \cup (\tilde{\Phi} \cap \tilde{\psi})) = 1$. Then by the preceding argument we have $\tilde{\Phi} \cap \tilde{\phi} \subseteq \tilde{\rho}$ and $\tilde{\Phi} \cap \tilde{\psi} \subseteq \tilde{\rho}$, which implies

$$(\tilde{\Phi} \cap \tilde{\phi}) \cup (\tilde{\Phi} \cap \tilde{\psi}) \subseteq \tilde{\rho}.$$

It follows that

$$1 = \nu((\tilde{\Phi} \cap \tilde{\phi}) \cup (\tilde{\Phi} \cap \tilde{\psi})) \leq \nu(\tilde{\rho});$$

whence, $\nu(\tilde{\rho}) = 1$ and $(\Omega, \mathcal{F}, \nu) \in \mu_{\rho}$, as claimed.

We next discuss the **0** type.

- **0-type**, formation rule:

Syntax	Semantics
$\Delta \text{ ctx}$	$\Delta \subseteq \mathcal{P}(\Pi^*)$
<hr style="width: 100%;"/>	<hr style="width: 100%;"/>
$\Delta \vdash \mathbf{0} : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_{\perp}$

By assumption Δ is a subset of probability spaces $(\Omega, \mathcal{F}, \nu)$, with \mathcal{F} a σ -algebra over Ω . As any σ -algebra contains \emptyset it is immediate that $\mathcal{P}(\Pi^*) = \mathcal{S}_{\perp}$ and the conclusion follows.

- **0 type**, elimination rule:²⁹

Syntax	Semantics
$\Delta \vdash \phi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_{\phi}$
<hr style="width: 100%;"/>	<hr style="width: 100%;"/>
$\Delta, x : \mathbf{0} \vdash \mathbf{efq}_{\phi}(x) : \phi$	$\Delta \cap \mu_{\perp} \subseteq \mu_{\phi}$

The subset $\mu_{\perp} \subseteq \mathcal{P}(\Pi^*)$ consists of all probability spaces $(\Omega, \mathcal{F}, \nu)$ that assign measure 1 to \emptyset . By definition, any probability measure ν must satisfy $\nu(\emptyset) = 0$, so that $\mu_{\perp} = \emptyset$. Thus, the conclusion reads $\Delta \cap \mu_{\perp} \equiv \emptyset \subseteq \mu_{\phi}$, which holds trivially.

²⁹Here **efq** stands for *ex falso quodlibet* (“from falsehood, anything follows”). Formally, this rule says that given any $\phi : \mathbf{Type}$ and a proof $x : \mathbf{0}$ of the contradiction it is possible to construct a proof $\mathbf{efq}_{\phi}(x) : \phi$. This is a type-theoretic version of the principle of explosion in PL, $\perp \rightarrow \phi$.

Finally, for the **Typ**-type.

- **Typicality type**, formation rule:

Syntax	Semantics
$\Delta \vdash \phi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\phi$
<hr style="width: 100%;"/>	<hr style="width: 100%;"/>
$\Delta \vdash \mathbf{Typ}(\phi) : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_{\mathbf{Typ}(\phi)}$

By definition, we require that $\widetilde{\mathbf{Typ}(\phi)}$ is a measurable set whenever $\tilde{\phi}$ is, so that the conclusion immediately follows by definition of $\mathcal{P}(\Pi^*)$.

- **Typicality type**, introduction rule:

Syntax	Semantics
$\Delta \vdash \phi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\phi$
<hr style="width: 100%;"/>	<hr style="width: 100%;"/>
$\Delta, a : \phi \vdash \tau_\phi(a) : \mathbf{Typ}(\phi)$	$\Delta \cap \mu_\phi \subseteq \mu_{\mathbf{Typ}(\phi)}$

As any $(\Omega, \mathcal{F}, \nu) \in \Delta \cap \mu_\phi$ must satisfy $\nu(\tilde{\phi}) = 1$ and it is required that $\widetilde{\mathbf{Typ}(\phi)} \supseteq \tilde{\phi}$, we must have $\nu(\widetilde{\mathbf{Typ}(\phi)}) \geq \nu(\tilde{\phi}) = 1$, and $(\Omega, \mathcal{F}, \nu) \in \mu_{\mathbf{Typ}(\phi)}$, as claimed.

- **Typicality type**, elimination rule:

Syntax	Semantics
$\Delta \vdash \phi : \mathbf{Type} \quad \Delta \vdash \psi : \mathbf{Type}$	$\Delta \subseteq \mathcal{S}_\phi \quad \Delta \subseteq \mathcal{S}_\psi$
$\Delta, a : \phi \vdash d(a) : \psi$	$\Delta \cap \mu_\phi \subseteq \mu_\psi$
<hr style="width: 100%;"/>	<hr style="width: 100%;"/>
$\Delta, x : \mathbf{Typ}(\phi) \vdash \mathbf{imp}_d(x) : \mathbf{Typ}(\psi)$	$\Delta \cap \mu_{\mathbf{Typ}(\phi)} \subseteq \mu_{\mathbf{Typ}(\psi)}$

By (3), we can express the second assumption as $\mu_{\Phi \wedge \phi} \subseteq \mu_\psi$ for some $\Phi \in \Pi^*$, from which it follows that $\tilde{\Phi} \cap \tilde{\phi} \subseteq \tilde{\psi}$ by an argument already given above when proving the elimination rule for the coproduct type. Thus, we can rewrite the conclusion as

$$\mu_{\Phi \wedge \mathbf{Typ}(\phi)} \subseteq \mu_{\mathbf{Typ}(\psi)}.$$

Let $(\Omega, \mathcal{F}, \nu) \in \mu_{\Phi \wedge \text{Typ}(\phi)}$. Then $\nu(\tilde{\Phi} \cap \widetilde{\text{Typ}(\phi)}) = 1$, and in particular $\nu(\tilde{\Phi}) = 1$ and $\nu(\widetilde{\text{Typ}(\phi)}) = 1$, implying that $\nu(\tilde{\phi}) \geq 1 - \epsilon$.

By definition we have

$$\begin{aligned} \nu(\tilde{\Phi} \cap \tilde{\phi}) &= \nu(\tilde{\Phi}) + \nu(\tilde{\phi}) - \nu(\tilde{\Phi} \cup \tilde{\phi}) \\ &\geq \nu(\tilde{\Phi}) + (1 - \epsilon) - 1 \\ &= 1 - \epsilon. \end{aligned}$$

By assumption, we have $\nu(\tilde{\psi}) \geq \nu(\tilde{\Phi} \cap \tilde{\phi}) \geq 1 - \epsilon$, and $\nu(\widetilde{\text{Typ}(\psi)}) = 1$ by definition.

• **Typicality type, 0 rule:**

Syntax	Semantics
$\Delta \text{ ctx}$	$\Delta \subseteq \mathcal{P}(\Pi^*)$
<hr style="width: 100%;"/>	<hr style="width: 100%;"/>
$\Delta, x : \mathbf{Typ}(\mathbf{0}) \vdash \sigma(x) : \mathbf{0}$	$\Delta \cap \mu_{\text{Typ}(\perp)} \subseteq \mu_{\perp}$

By definition, every probability measure is required to assign probability 0 to $\tilde{\perp} \equiv \emptyset$. Thus there does not exist any probability measure ν for which $\nu(\emptyset) \geq 1 - \epsilon$, as is required to make the judgment that \perp is typical. It follows that $\mu_{\text{Typ}(\perp)} \equiv \emptyset$; whence, $\Delta \cap \mu_{\text{Typ}(\perp)} \equiv \emptyset \subseteq \emptyset \equiv \mu_{\perp}$, as required.

This completes the proof of soundness.

□

3.5 Additional Results

The above rules for the **Typ**-type have several immediate consequences for derived inference rules involving typicality. For example, we have the following one-way hierarchy of typicality judgments involving conjunction and disjunction. With ‘ \Rightarrow ’ understood informally as ‘implies’, we have

$$\text{Typ}(\phi \wedge \psi) \Rightarrow \text{Typ}(\phi) \wedge \text{Typ}(\psi) \Rightarrow \text{Typ}(\phi) \vee \text{Typ}(\psi) \Rightarrow \text{Typ}(\phi \vee \psi). \quad (4)$$

We prove this formally below.

Theorem 7. *The implication arrows in (4) correspond, respectively, to the following derived inference rules:³⁰*

1. $\text{Typ}(\phi \wedge \psi) \Rightarrow \text{Typ}(\phi) \wedge \text{Typ}(\psi)$:

$$\frac{\Delta \vdash \phi : \mathbf{Type} \quad \Delta \vdash \psi : \mathbf{Type}}{\Delta, z : \mathbf{Typ}(\phi \times \psi) \vdash (\text{imp}_{\text{pr}_\phi}(z), \text{imp}_{\text{pr}_\psi}(z)) : \mathbf{Typ}(\phi) \times \mathbf{Typ}(\psi)}$$

2. $\text{Typ}(\phi) \wedge \text{Typ}(\psi) \Rightarrow \text{Typ}(\phi) \vee \text{Typ}(\psi)$.³¹

$$\frac{\Delta \vdash \phi : \mathbf{Type} \quad \Delta \vdash \psi : \mathbf{Type}}{\Delta, z : \mathbf{Typ}(\phi) \times \mathbf{Typ}(\psi) \vdash \text{inl}(\text{pr}_{\mathbf{Typ}(\phi)}(z)) : \mathbf{Typ}(\phi) + \mathbf{Typ}(\psi)}$$

3. $\text{Typ}(\phi) \vee \text{Typ}(\psi) \Rightarrow \text{Typ}(\phi \vee \psi)$:

$$\frac{\Delta \vdash \phi : \mathbf{Type} \quad \Delta \vdash \psi : \mathbf{Type}}{\Delta, z : \mathbf{Typ}(\phi) + \mathbf{Typ}(\psi) \vdash \text{case}_{\text{imp}_{\text{inl}}, \text{imp}_{\text{inr}}}(z) : \mathbf{Typ}(\phi \vee \psi)}$$

Proof. To prove the first implication, we observe first that

$$\Delta, x : \phi \times \psi \vdash \text{pr}_\phi(x) : \phi \quad \text{and}$$

$$\Delta, x : \phi \times \psi \vdash \text{pr}_\psi(x) : \psi$$

are both valid judgments in MLTT which match the second line in the elimination rule for

³⁰Note below we define the projection operator pr_ϕ as that which selects out the corresponding coordinate from a pair $(a, b) : \phi \times \psi$. For example, $\text{pr}_\phi((a, b)) \equiv a : \phi$ and $\text{pr}_\psi((a, b)) \equiv b : \psi$.

³¹An alternative formalization of this implication is given by

$$\frac{\Delta \vdash \phi : \mathbf{Type} \quad \Delta \vdash \psi : \mathbf{Type}}{\Delta, z : \mathbf{Typ}(\phi) \times \mathbf{Typ}(\psi) \vdash \text{inr}(\text{pr}_{\mathbf{Typ}(\psi)}(z)) : \mathbf{Typ}(\phi) + \mathbf{Typ}(\psi)}$$

the **Typ**-type. It follows, from each of these lines respectively, that

$$\Delta, y : \mathbf{Typ}(\phi \times \psi) \vdash \mathbf{imp}_{\mathbf{pr}_\phi}(y) : \mathbf{Typ}(\phi) \quad \text{and}$$

$$\Delta, y : \mathbf{Typ}(\phi \times \psi) \vdash \mathbf{imp}_{\mathbf{pr}_\psi}(y) : \mathbf{Typ}(\psi).$$

An application of the introduction rule for the product type gives

$$\Delta, y : \mathbf{Typ}(\phi \times \psi) \vdash (\mathbf{imp}_{\mathbf{pr}_\phi}(y), \mathbf{imp}_{\mathbf{pr}_\psi}(y)) : \mathbf{Typ}(\phi) \times \mathbf{Typ}(\psi).$$

To prove the second implication, it is enough to observe that $\Phi \times \Psi \Rightarrow \Phi + \Psi$ for all $\Phi, \Psi : \mathbf{Type}$ in any context. Formally, we have

$$\Delta, x : \mathbf{Typ}(\phi) \times \mathbf{Typ}(\psi) \vdash \mathbf{inl}(\mathbf{pr}_{\mathbf{Typ}(\phi)}(x)) : \mathbf{Typ}(\phi) + \mathbf{Typ}(\psi),$$

or alternatively

$$\Delta, x : \mathbf{Typ}(\phi) \times \mathbf{Typ}(\psi) \vdash \mathbf{inr}(\mathbf{pr}_{\mathbf{Typ}(\psi)}(x)) : \mathbf{Typ}(\phi) + \mathbf{Typ}(\psi).$$

To prove the third implication, we must apply the elimination rule for the co-product type. In this case, the elimination rule is applied by

$$\begin{array}{c} \Delta \vdash \mathbf{Typ}(\phi) : \mathbf{Type} \quad \Delta \vdash \mathbf{Typ}(\psi) : \mathbf{Type} \\ \Delta, a : \mathbf{Typ}(\phi), b : \mathbf{Typ}(\psi) \vdash \mathbf{Typ}(\phi + \psi) : \mathbf{Type} \\ \Delta, a : \mathbf{Typ}(\phi) \vdash \mathbf{imp}_{\mathbf{inl}}(a) : \mathbf{Typ}(\phi + \psi) \\ \Delta, b : \mathbf{Typ}(\psi) \vdash \mathbf{imp}_{\mathbf{inr}}(b) : \mathbf{Typ}(\phi + \psi) \\ \hline \Delta, z : \mathbf{Typ}(\phi) + \mathbf{Typ}(\psi) \vdash \mathbf{case}_{\mathbf{imp}_{\mathbf{inl}}, \mathbf{imp}_{\mathbf{inr}}}(z) : \mathbf{Typ}(\phi + \psi) \end{array}$$

□

We note that the arrows in (4) do not reverse in general. To see this, we can produce a semantic counterexample by giving a probability space $(\Omega, \mathcal{F}, \nu)$ for which the corresponding interpretation fails.

1. $Typ(\phi \wedge \psi) \not\equiv Typ(\phi) \wedge Typ(\psi)$. To show this, we instead show

$$\mu_{Typ(\phi \wedge \psi)} \not\equiv \mu_{Typ(\phi)} \cap \mu_{Typ(\psi)}.$$

Let $\Omega = \{\omega_1, \omega_2, \omega_3\}$ have the power set σ -algebra and define measurable sets corresponding to ϕ , ψ , $Typ(\phi)$, $Typ(\psi)$, and their conjunctions and disjunctions by

$$\begin{aligned}\tilde{\phi} &= \{\omega_1, \omega_2\} \\ \tilde{\psi} &= \{\omega_2, \omega_3\} \\ \widetilde{Typ(\phi)} &= \{\omega_1, \omega_2, \omega_3\} = \Omega \\ \widetilde{Typ(\psi)} &= \{\omega_1, \omega_2, \omega_3\} = \Omega \\ \widetilde{Typ(\phi \wedge \psi)} &= \{\omega_2\}\end{aligned}$$

Fix $\epsilon = 1/3$ and define measure ν by

$$\nu(\{\omega_1\}) = \nu(\{\omega_3\}) = 1/4 \quad \text{and} \quad \nu(\{\omega_2\}) = 1/2.$$

From this, we have

$$\nu(\tilde{\phi}) = \nu(\tilde{\psi}) = 3/4 \geq 1 - \epsilon,$$

so that ϕ and ψ are typical according to ν . (Also note that

$$\nu(\widetilde{Typ(\phi)}) = \nu(\widetilde{Typ(\psi)}) = \nu(\Omega) = 1,$$

as required to justify the judgments that ϕ and ψ are typical.) Thus, we have $Typ(\phi) \wedge Typ(\psi)$ in accordance with the righthand side, but

$$\nu(\tilde{\phi} \cap \tilde{\psi}) = \nu(\{\omega_2\}) = 1/2 < 1 - \epsilon,$$

so that $\phi \wedge \psi$ is not typical.

2. $Typ(\phi) \wedge Typ(\psi) \not\equiv Typ(\phi) \vee Typ(\psi)$. To show this, we instead show

$$\mu_{Typ(\phi)} \cap \mu_{Typ(\psi)} \not\equiv \mu_{Typ(\phi)} \cup \mu_{Typ(\psi)}.$$

Let $\Omega = \{\omega_1, \omega_2, \omega_3\}$ have the power set σ -algebra and define measurable sets corresponding to ϕ , ψ , $Typ(\phi)$, $Typ(\psi)$, and their conjunctions and disjunctions by

$$\begin{aligned}\tilde{\phi} &= \{\omega_1, \omega_2\} \\ \tilde{\psi} &= \{\omega_2, \omega_3\} \\ \widetilde{Typ(\phi)} &= \{\omega_1, \omega_2, \omega_3\} = \Omega \\ \widetilde{Typ(\psi)} &= \{\omega_2, \omega_3\} \\ \widetilde{Typ(\phi \wedge \psi)} &= \{\omega_2\}\end{aligned}$$

Fix $\epsilon = 1/3$ and define measure ν by

$$\nu(\{\omega_1\}) = 1 \quad \text{and} \quad \nu(\{\omega_2\}) = \nu(\{\omega_3\}) = 0.$$

Thus, $\nu(\tilde{\phi}) = 1$ justifies the judgment in $Typ(\phi)$, so that the righthand side holds. But $\widetilde{Typ(\phi)} \cap \widetilde{Typ(\psi)} = \{\omega_2, \omega_3\}$ has measure 0 under ν , so the lefthand side does not hold under ν .

3. $Typ(\phi) \vee Typ(\psi) \not\equiv Typ(\phi \vee \psi)$. To show this, we instead show

$$\mu_{Typ(\phi)} \cup \mu_{Typ(\psi)} \not\equiv \mu_{Typ(\phi \vee \psi)}.$$

Let $\Omega = \{\omega_1, \omega_2, \omega_3\}$ have the power set σ -algebra and define measurable sets corresponding to ϕ , ψ , $Typ(\phi)$, $Typ(\psi)$, and their conjunctions and disjunctions by

$$\tilde{\phi} = \{\omega_1, \omega_2\}$$

$$\begin{aligned}
\tilde{\psi} &= \{\omega_2, \omega_3\} \\
\widetilde{Typ}(\phi) &= \{\omega_1, \omega_2\} \\
\widetilde{Typ}(\psi) &= \{\omega_2, \omega_3\} \\
\widetilde{Typ}(\phi \wedge \psi) &= \{\omega_2\} \\
\widetilde{Typ}(\phi \vee \psi) &= \{\omega_1, \omega_2, \omega_3\} = \Omega
\end{aligned}$$

Fix $\epsilon = 1/10$ and define measure ν by

$$\nu(\{\omega_1\}) = \nu(\{\omega_3\}) = 1/4 \quad \text{and} \quad \nu(\{\omega_2\}) = 1/2.$$

In this case, we have

$$\nu(\tilde{\phi}) = \nu(\tilde{\psi}) = 0.75 < 1 - 1/10,$$

so that neither ϕ nor ψ is typical and the lefthand side fails to hold. But the disjunction $\tilde{\phi} \cup \tilde{\psi} = \Omega$ has ν -measure 1, and is therefore typical, satisfying the righthand side.

4 Conclusion

Formally, TPL and TIL have a great deal in common. Their proof theories are sound with respect to their semantic theories. In both TPL and TIL, typicality distributes over conjunction in one direction but not in the other: if $p \wedge q$ is typical then p is typical and q is typical, but if p is typical and q is typical then it does not follow that $p \wedge q$ is typical. Similarly, in both TPL and TIL, typicality distributes over disjunction in one direction but not in the other: if p is typical or q is typical then $p \vee q$ is typical, but if $p \vee q$ is typical then it does not follow that p is typical or q is typical.

There are important formal differences between TPL and TIL, however. The logic underlying TPL is classical, whereas the logic underlying TIL is intuitionistic. And while

TPL is formulated in the manner of propositional modal logic, TIL is formulated in the manner of type theory.

These formal differences correspond to conceptual and metaphysical differences between the notion of typicality captured by TPL and the notion of typicality captured by TIL. For instance, according to TPL, the notion of typicality conforms to the rules of classical logic. According to TIL, in contrast, the notion of typicality is more constructive than that, requiring that an agent make explicit judgments about the truth and typicality of propositions. And there are other conceptual and metaphysical differences between these two systems. Whereas the meanings of typicality statements in TPL are given by sets of possible worlds, the meanings of typicality statements in TIL are given by sets of possible credences. So TPL is perhaps best understood as formalizing objective typicality facts, while TIL is perhaps best understood as formalizing subjective typicality judgments. Given these differences, it is worth exploring – in future work – the circumstances in which one system may be preferable to the other.

But regardless, both systems offer rigorous regimentations of typicality reasoning. They formalize the logical structure of the notion of typicality: its grammar, its semantic content, its proof theory, and the sorts of valid inferences which it licenses. Because of that, TPL and TIL offer formal frameworks for typicality reasoning on a par with the formal framework that first-order logic offers for quantificational reasoning, or the formal framework that Bayesian theory offers for reasoning in terms of credences. TPL and TIL limn the deep logical structure which is shared by many instances of typicality reasoning in the literature on quantum mechanics and statistical mechanics. These systems reveal, in other words, what is common to the diverse array of explanations, predictions, and other kinds of scientific reasoning, which invoke the notion of the typical. In short, TPL and TIL capture the logic of typicality.

References

- Alon, N., Krivelevich, M., & Sudakov, B. (1998). Finding a Large Hidden Clique in a Random Graph. *Random Structures & Algorithms*, 13, (3-4), 457-466.
- Bernoulli, J. (1713). *Ars Conjectandi*. Thurnisius, Basel.
- Boltzmann, L. (2003). Reply to Zermelo's Remarks on the Theory of Heat. In S. G. Brush (Ed.), *The Kinetic Theory of Gases* (pp. 392-402). London: Imperial College Press. (Original work published 1896)
- Booth, R., Meyer, T., & Varzinczak, I. (2012). PTL: A Propositional Typicality Logic. In L. F. del Cerro, A. Herzig, & J. Mengin (Eds.), *Logics in Artificial Intelligence* (pp. 107-119), Heidelberg: Springer.
- Cournot, A. A. (1843). *Exposition de la théorie des chances et des probabilités*. Hachette, Paris.
- Crane, H. (2018). Logic of probability and conjecture. *Researchers.One*. Accessed on February 5, 2019 at <https://www.researchers.one/article/2018-08-5>.
- Curry, H. & Feys, R. (1958). *Combinatory Logic. I*. Amsterdam: North-Holland Publishing Company.
- Dürr, D., Goldstein, S., & Zanghì, N. (1992). Quantum Equilibrium and the Origin of Absolute Uncertainty. *Journal of Statistical Physics*, 67(5/6), 843-907.
- Everett, H. (2012). The Theory of the Universal Wave Function. In J. Barrett & P. Byrne (Eds.), *The Everett Interpretation of Quantum Mechanics* (pp. 72-172). Princeton, NJ: Princeton University Press. (Original work published 1956)
- Frigg, R. (2011). Why Typicality Does Not Explain the Approach to Equilibrium. In M. Suárez (Ed.), *Probabilities, Causes and Propensities in Physics* (pp. 77-93). New York, NY: Springer.
- Frigg, R., & Werndl, C. (2012). Demystifying Typicality. *Philosophy of Science*, 79, 917-929.
- Goldstein, S. (2001). Boltzmann's Approach to Statistical Mechanics. In J. Bricmont, D. Dürr, M. C. Galavotti, G. Ghirardi, F. Petruccione, & N. Zanghì (Eds.), *Chance in*

- Physics* (pp. 39-54). Heidelberg: Springer.
- Goldstein, S. (2012). Typicality and Notions of Probability in Physics. In Y. Ben-Menahem & M. Hemmo (Eds.), *Probability in Physics* (pp. 59-71). New York, NY: Springer-Verlag.
- Goldstein, S., Lebowitz, J., Mastrodonato, C., Tumulka, R., & Zanghì, N. (2010). Normal typicality and von Neumann's quantum ergodic theorem. *Proceedings: Mathematical, Physical and Engineering Sciences*, *466*(2123), 3203-3224.
- Goldstein, S., Lebowitz, J. L., Tumulka, R., & Zanghì, N. (2006). Canonical Typicality. *Physical Review Letters*, *96*, 050403.
- Howard, W. A. (1969). The formulae-as-types notion of construction. In J. P. Seldin, R. J. Hindley, & H. B. Curry: *Essays on Combinatory Logic, lambda calculus and formalism* (pp. 479-490). New York, NY: Academic Press.
- Kesten, H. (1980). The Critical Probability of Bond Percolation on the Square Lattice Equals $1/2$. *Communications in Mathematical Physics*, *74*, 41-59.
- Kiessling, M. K-H. (2011). Typicality analysis for the Newtonian N -body problem on \mathbb{S}^2 in the $N \rightarrow \infty$ limit. *Journal of Statistical Mechanics: Theory and Experiment*, P01028.
- Ledoux, M. (2001). *The Concentration of Measure Phenomenon*. Providence, RI: American Mathematical Society.
- Martin-Löf, P. (1984). *Intuitionistic type theory*. Sambin, Giovanni. Napoli: Bibliopolis.
- Mazur, B. (2012). Is it plausible? Available at <http://www.math.harvard.edu/mazur/papers/Plausibility.Notes.3.pdf>. Accessed on August 20, 2018.
- Priest, G. (2008). *An Introduction to Non-Classical Logic*. New York, NY: Cambridge University Press.
- Reimann, P. (2007). Typicality for Generalized Microcanonical Ensembles. *Physical Review Letters*, *99*, 160404.
- Shafer, G., & Vovk, V. (2006). The Sources of Kolmogorov's "Grundbegriffe". *Statistical Science*, *21*(1), 70-98.

- Smullyan, R. M. (1968). *First-Order Logic*. New York, NY: Springer-Verlag.
- Tasaki, H. (2016). Typicality of Thermal Equilibrium and Thermalization in Isolated Macroscopic Quantum Systems. *Journal of Statistical Physics*, 163, 937-997.
- Tsementzis, D. (2018). A Meaning Explanation for HoTT. Available at <http://rci.rutgers.edu/dt506/A.Meaning.Explanation.for.HoTT.v3.5.public.pdf>. Accessed on August 27, 2018.
- Werndl, C. (2013). Justifying typicality measures of Boltzmannian statistical mechanics and dynamical systems. *Studies in History and Philosophy of Modern Physics*, 44, 470-479.
- Wilhelm, I. (forthcoming). Typical: A Theory of Typicality and Typicality Explanation. *The British Journal for the Philosophy of Science*.