

Understanding Universal Disjunction

Garry Goodwin

January 11, 2026

Abstract

The difference in meaning between the two versions of universal disjunction is rarely subjected to systematic analysis. We identify and examine twelve distinct interpretations that allow us to distinguish individually quantified disjoint sentences from cases in which the universal quantifier is distributed over a disjunction.

In quantificational logic (QL), *universal disjunction* arises in two formally distinct ways. On the one hand, one may form a disjunction of universally quantified sentences, for example,

$$\forall x(Bx) \vee \forall x(Rx).$$

On the other, one may place a universal quantifier in front of a disjunction, thereby allowing the quantifier to distribute over the disjunctive matrix,

$$\forall x(Bx \vee Rx).$$

Lemmon illustrates the logical difference between these two constructions with a well-known example.^[1]

Let the domain consist of the positive integers. From the fact that every number is either even or odd, it does *not* follow that all numbers are even, nor that all numbers are odd. This example makes it clear that entailment holds in one direction but fails in the other. As Equation (1) makes explicit, this asymmetry is a basic tenet of QL.

$$\forall x(Bx \vee Rx) \not\vdash \forall x(Bx) \vee \forall x(Rx). \quad (1)$$

While this result is provable within the syntactic rules of QL, and the positive integer example provides a clear-cut case, sustained attempts to give a systematic explanation of why the meanings of such sentences do not permit reverse entailment are rare. As we shall see, Lemmon’s example leaves much out. To show what is missing, our investigation is confined to the following stock of symbols, governed by the usual formation rules for well-formed formulae:

$$\forall, \exists, B, R, x, a, \neg, \wedge, \vee, \rightarrow, (,).$$

Here B and R are unary predicate letters, a is a constant symbol, and x is the sole variable. Although this restricted vocabulary is still able to generate infinitely many well-formed formulae, the majority are logically equivalent. Indeed, only a finite number of logically distinct meanings can be expressed within this limited fragment of QL.

To account for this finite set, it is helpful to introduce the concept of a *semantic atom*—that is, a minimal satisfiable proposition within the fragment. Unlike an atomic sentence, a semantic atom may be syntactically complex, while nevertheless corresponding to a single proposition that is entailed only by sentences that are logically equivalent to it, or by sentences that are logically contradictory. Two sentences express the same meaning precisely when the standard derivation rules prove that they mutually entail one another; a semantic atom is therefore represented by an equivalence class of sentences.

What remains to be explained, in order to distinguish the two forms of universal disjunction, is a semantics capable of showing when this mutual entailment holds and when it does not. The present method relies on the fact that, in this restricted fragment, there are only finitely many equivalence classes of sentences; in richer languages, this kind of finite classification need not be available. Since logical equivalence is not—in general—decidable, the explicit identification and enumeration of semantic atoms is most straightforwardly limited to fragments in which equivalence is decidable.

Universal disjunctions in which the quantifier is not distributed may then be distinguished with a 4×4 array of Boolean ones and zeroes. However, it will be helpful to represent an array graphically with each Boolean value shown as a “tile.” A white (ivory) tile is a Boolean 1 (true) and a black tile is a Boolean 0 (false). The sixteen tiles therefore allow 2^{16} distinct Boolean patterns, each corresponding to a possible combination of the four propositions illustrated in Figure 1 (namely Ba , Ra , $\forall x(Bx)$, and $\forall x(Rx)$).




Figure 1

Unfortunately, the basic 4×4 array is insufficient to express distributed universal disjunction; for this we need the extended array of Figure 2. The full set of thirty-two tiles gives rise to 2^{32} distinct combinations of meanings formed from these tiles, a number approaching 4.3 billion; large, but still finite.




Figure 2

With the four central tiles now each replaced with a system of five tiles, the double triangle (hourglass) represents a single proposition. Although its precise shape is arbitrary, as we shall see later with P11 and P12, the double triangle invokes two possibilities that exclude one another.

Owing to the obvious resemblance, we refer to such grid-based representations as *mosaics*. The mosaic pattern *points* to the respective corner of the array as shown in Figure 3.




Figure 3

To illustrate the difference in meaning between the two forms of universal disjunction, we restrict attention to the positive predicates

represented by the two forms of universal disjunction mosaics shown in Figure 4.




Figure 4

Mosaics are not themselves semantic objects, but graphical representations of semantic distinctions; as such they are essentially a bookkeeping device for tracking entailment relations. Since each tile corresponds to an infinite number of well-formed formulae, we therefore select, in order to keep the discussion manageable, a single formula for each proposition, chosen to have the simplest available syntax. The domain we shall consider consists of flower stalks, each bearing blue (B) or red (R) petals.


The following analysis yields twelve distinct meanings, labelled P1 to P12. Each is illustrated by an English sentence (or sentences), an accompanying picture, and a Venn diagram.

P1. $\forall x(Bx) \wedge \forall x(Rx)$

Every stalk has a blue petal and a red petal.


$$P2. \forall x(Bx) \wedge \forall x(\neg Rx)$$

Every stalk has a blue petal, none has a red petal.




$$P3. \forall x(Rx) \wedge \forall x(\neg Bx)$$

Every stalk has a red petal, none has a blue petal.




$$P4. \forall x(Bx) \wedge \exists x(\neg Rx) \wedge Ra$$

Every stalk has a blue petal.
Some stalks do not have a red petal, but stalk *a* does.




P5. $\forall x(Rx) \wedge \exists x(\neg Bx) \wedge Ba$

Every stalk has a red petal.
Some stalks do not have a blue petal, but stalk a does.




P6. $\forall x(Bx) \wedge \exists x(Rx) \wedge \neg Ra$

Every stalk has a blue petal.
Some stalks have a red petal, but stalk a does not.




P7. $\forall x(Rx) \wedge \exists x(Bx) \wedge \neg Ba$

Every stalk has a red petal.
Some stalks have a blue petal, but stalk a does not.




$$P8. \forall x(Bx \vee Rx) \wedge \exists x(\neg Bx) \wedge \exists x(\neg Rx) \wedge Ba \wedge Ra$$

Every stalk has a blue or red petal.
 Some stalks do not have a blue petal
 and some do not have red, but stalk a has both.




$$P9. \forall x(Bx \vee Rx) \wedge \exists x(Bx \wedge Rx) \wedge \exists x(\neg Bx) \wedge \neg Ra$$

Every stalk has a blue petal or red petal.
 Some stalks have both a blue petal and a red petal.
 Some stalks do not have a blue petal.
 Stalk a does not have a red petal.




$$P10. \forall x(Bx \vee Rx) \wedge \exists x(Bx \wedge Rx) \wedge \exists x(\neg Rx) \wedge \neg Ba$$

Every stalk has a blue or red petal.
 Some stalks have both a blue petal and a red petal.
 Some stalks do not have a red petal.
 Stalk a does not have a blue petal.




P11. $\forall x(Bx \rightarrow \neg Rx) \wedge \forall x(\neg Bx \rightarrow Rx) \wedge \exists x(\neg Bx) \wedge \exists x(Rx) \wedge \neg Ra$

If stalks have a blue petal then they do not have a red petal and if stalks do not have a blue petal they have a red petal.
 Some stalks do not have a blue petal.
 Some stalks have a red petal, but stalk a does not.



P12. $\forall x(Bx \rightarrow \neg Rx) \wedge \forall x(\neg Bx \rightarrow Rx) \wedge \exists x(\neg Rx) \wedge \exists x(Bx) \wedge \neg Ba$

If stalks have a blue petal then they do not have a red petal and if they do not have a blue petal they have a red petal. Some stalks do not have a red petal.
 Some stalks have a blue petal, but stalk a does not.



Perhaps the most surprising result of this analysis is that, to characterise universal quantification within this fragment, it is necessary to name at least one object of the domain in nine of the twelve propositions. Across the series, the named element is explicitly overlaid on the corresponding Venn diagrams: depending on the proposition in question, it may appear in the B -region, the R -region, or their intersection. However, whether the associated closed atomic sentence is positive or negative is not determined by this placement alone and only becomes evident with the inclusion of the array. While the pictorial representation of the flowers provides additional illumination, it is not strictly necessary; the array and the Venn diagram should

therefore be understood together as a single semantic framework, capable of representing the distinct relationships exhibited by this set of formulae.

Admittedly, we have simply stated the set of twelve propositions without explicit derivation. Nevertheless, the accompanying Venn diagrams provide assurance that all logical possibilities have been exhaustively represented. Moreover, careful study will confirm that this analysis is supported by four arguments valid in QL. We omit the formal proofs, which are routine but lengthy.

The first establishes that any two propositions of the set P1–P12 cannot both be true simultaneously.

$$\vdash \neg(Pn \wedge Pm),$$

where Pn and Pm are any two propositions taken from P1 to P12.

This theorem is needed if the twelve propositions are semantic atoms. Although the proof is not reproduced here, careful attention to the meaning of the sentences themselves will show that this is indeed a set of contrary statements.

The next valid argument confirms the disjunction of two universally quantified sentences is equal to the disjunction P1 to P7.

$$\forall x(Bx) \vee \forall x(Rx) \dashv\vdash \bigvee (P1 - P7).$$

Given that each of P1 through P7 is a minimal statement (in the sense that it is a semantic atom implied only by a logically equivalent sentence or by a contradiction), and that any two of these propositions form a contrary pair, their full disjunction provides a natural semantic normal form for $\forall x(Bx) \vee \forall x(Rx)$.

The third valid argument confirms distributed universal quantification is equal to the disjunction P1 to P12.

$$\forall x(Bx \vee Rx) \dashv\vdash \bigvee (P1 - P12).$$

There is also a theorem for the mosaic of thirty-two tiles confirming the mosaic with thirty-two tiles introduced in Figure 2 exhausts all possible relations for this QL fragment.

$$\vdash \bigvee (P1 - P32).$$

A disjunction of thirty-two propositions would clearly constitute a cumbersome derivation. However, the symmetry of the mosaic means

that it suffices to show that the eight propositions forming the top-right quartile are jointly equivalent to $Ba \wedge Ra$. The eight propositions are:

$$P1 : \forall x(Bx) \wedge \forall x(Rx)$$

$$P4 : \forall x(Bx) \wedge \exists x(\neg Rx) \wedge Ra$$

$$P5 : \forall x(Rx) \wedge \exists x(\neg Bx) \wedge Ba$$

$$P8 : Ba \wedge Ra \wedge \forall x(Bx \vee Rx) \wedge \exists x(\neg Bx) \wedge \exists x(\neg Rx)$$

$$P13 : Ba \wedge Ra \wedge \forall x(\neg Bx \vee Rx) \wedge \forall x(Bx \vee \neg Rx) \wedge \exists x(\neg Bx) \wedge \exists x(\neg Rx)$$

$$P14 : Ba \wedge Ra \wedge \neg(P13) \wedge \forall x(\neg Bx \vee Rx)$$

$$P15 : Ba \wedge Ra \wedge \neg(P13) \wedge \forall x(Bx \vee \neg Rx)$$

$$P16 : Ba \wedge Ra \wedge \exists x(\neg Bx \wedge Rx) \wedge \exists x(Bx \wedge \neg Rx) \wedge \exists x(\neg Bx \wedge \neg Rx)$$

As before, each of these propositions is a semantic atom and any two form a contrary pair. Thus,

$$Ba \wedge Ra \dashv\vdash \bigvee (P1, P4, P5, P8, P13, P14, P15, P16).$$

It is worth pausing to consider what this result shows. Within this fragment of quantificational logic, the simple conjunction $Ba \wedge Ra$, formed from closed atomic formulae, can nevertheless correspond to a disjunction whose disjuncts each contain quantified structure. Even sentences that are syntactically quantifier-free may therefore encode genuinely non-redundant quantificational content at the level of meaning, despite the absence of any overt quantificational syntax. The use of closed atomic formulae alone does not, therefore, preclude quantificational structure at the semantic level.

In conclusion, the mosaic representations help to illuminate the difference between the two forms of universal disjunction by making explicit distinctions in meaning that are not brought out by standard examples. The advantages of this approach become most apparent when Figure 4 serves as a reminder of why Equation (1) holds.

All theorems and valid forms of argument stated in this paper have been verified using the automated theorem prover Prover9.

References

- [1] Edward John Lemmon. *Beginning logic*. CRC Press, 1971.