An Efficient Communication Analysis of Modal Typology*

Nathaniel Imel
Department of Language Science
University of California, Irvine

Qingxia Guo
Department of Linguistics
University of Washington

Shane Steinert-Threlkeld
Department of Linguistics
University of Washington

Abstract The meanings expressed by the world's languages have been argued to support efficient communication. Across diverse semantic domains, crosslinguistic analyses show that natural language vocabularies are jointly optimized for two competing pressures: cognitive simplicity and informative communication. This paper applies an efficiency analysis to *modals* (e.g. *can, ought, might*). We define and measure the simplicity and informativeness of a large number of logically possible modal systems, including a sample of twenty seven natural language inventories. We also consider a recently-introduced semantic universal for modal expressions in natural language, dubbed the Independence of Force and Flavor (IFF). Our analysis yields three main results: (i) every optimal modal system perfectly satisfies the IFF universal; (ii) as systems contain more IFF modals, they become more efficient; (iii) attested modal systems are more efficient than merely possible systems. These results indicate that general pressures for efficient communication can explain typological variation in the lexicalization of modality.

Keywords: modals, typology, semantic universals, efficient communication

1 Introduction

The languages of the world exhibit *constrained variation*. While they differ substantially in important ways, there are also many possible but unattested languages, and actual languages exhibit considerable shared structure. Put differently, only a small subset of the mathematically possible languages have ever been spoken by

* We are grateful to Toshi Ogihara, Jakub Szymanik, Wataru Uegaki, Noga Zaslavsky, and audiences at Amsterdam, Austin, Davis, Irvine, and MIT for helpful discussion. We also thank James Engels, Ella Hannon, Anne Mucha, Wataru Uegaki, Fred Whibley, Minghe Zhang, Zoe Fang and Ian Chandra Lim for their work on contributing primary modals data to the modal typology database, described in Section 5.1

any linguistic community. One goal of theoretical linguistics consists in accurately characterizing and explaining this subset, i.e. in identifying the 'humanly possible' languages. Such explanations can occur at every level of linguistic analysis. The field of *semantic typology* asks: which meanings are attested cross-linguistically, and why (Bach et al. 2009, Kemp et al. 2018)?

In many domains, robust constraints on the meanings expressed in the languages of the world—*semantic universals*—have been discovered (Barwise et al. 1981, von Fintel & Matthewson 2008, Bach et al. 2009). When such universals are found, it is natural to want to explain them as well. An idea with roots in the functionalist tradition of linguistics proposes that the meanings observed across languages are shaped by a pressure for efficient communication (Zipf 1949, Kemp et al. 2018).

This paper evaluates this efficiency hypothesis in terms of the trade-off between cognitive simplicity and communicative informativeness (Kemp et al. 2012, Levinson 2012, Kemp et al. 2018), in the domain of *modals*, from the perspective both of semantic universals and typology more broadly. A modal is typically considered to be a semantic operator that qualifies the truth of an expression. In English, these can be expressed by auxiliaries including *might, may, must, could* and adverbs like *probably, necessarily* among a variety of other constructions. Cross-linguistically, these meanings are expressed by diverse lexical categories and strategies (Bhatt et al. 2011, Cable 2017).

Modality exemplifies the property that Hockett (1960) named displacement: the phenomenon of talking about beyond the actual here and now. Modals are typically also context-sensitive: words like *can* and *must* do not fully specify what kind of modality (i.e. relevant facts about knowledge, norms, or ability) a speaker has in mind, which means that listeners must rely on context to disambiguate the intended sense of possibility (see Section 2). In this paper, we explore whether efficiency can explain *why* modals lexicalize the specific combinations of possibility and necessity that are crosslinguistically attested. In other words, we are interested in whether efficiency can explain why words like *can* and *must* mean what they actually mean.

To evaluate the efficient communication hypothesis, we measure the complexity and informativeness achieved by a sample of 27 natural language modal systems, comparing them to a large number of hypothetical modal systems. We also consider a lexeme-level semantic universal that has been proposed to account for robust constraints on crosslinguistic variation in modal semantic systems, the Independence of Force and Flavor (IFF) (Steinert-Threlkeld et al. 2023). We observe that (i) every optimal modal system consists only of IFF modals; (ii) as systems contain more IFF modals, they become more efficient; (iii) attested modal systems are more efficient than merely possible systems. Together, these results suggest that communicative efficiency can provide a robust explanation both of the presence of a semantic universal and broader patterns in semantic typology in the domain of modals.

The paper is structured as follows. We first provide an overview of cross-linguistic semantic variation for modals (Section 2). We explain the IFF semantic universal in Section 3. We then introduce the simplicity/informativeness trade-off and describe in detail how to measure these properties in the modal domain (Section 4). The main computational experiment and results are presented in Section 5. After showing that all of the main results are robust to various alternative ways of measuring informativeness (Section 6), we conclude with a discusssion of modeling decisions and areas for future work (Section 7).

2 Modality

Modals are expressions that are used to talk about alternative ways the world could be, over and above the way the world actually is. Paradigms are certain English auxiliaries like *may* and *must*. Since at least Kratzer (1981), the semantics of modals have been explicated in terms of two axes of variation: **force** and **flavor**. These axes can be illustrated with the following examples.

(1) a. [Context: a friend walks in and shakes off a wet umbrella. You say:]
It *must* be raining.
b. [Context: you are reading the specifications of a homework assignment. It partially reads:]
You *must* upload your homework as a PDF.

(2) a. [Context: a friend is leaving and grabs an umbrella on the way out, saying:]
It *may* be raining.
b. [Context: a mother offers a treat to a child for finishing an assignment:]
You *may* have a cookie.

The *must* examples exhibit strong (i.e. universal) force, but differ in flavor. For example, (1a) can be glossed as saying: all of the worlds compatible with my evidence are worlds in which it is raining. The universal quantification represents the force, and the domain of worlds (those compatible with my evidence) the flavor, in this case epistemic. (1b) exhibits universal force with deontic flavor, roughly saying that all the worlds in which you follow the rules are ones in which you upload a PDF. The examples with *may* in (2) exhibit weak (i.e. possibility) force: their meaning says that some world satisfies the prejacent. (2a) and (2b) again differ in flavor, with the former being epistemic and the latter being deontic.

In addition to epistemic and deontic flavors, many others have been identified: bouletic (worlds in which desire are fulfilled), teleological (worlds in which goals are satisfied), et cetera. Similarly, there are arguably more forces than just weak and strong: for instance, there are weak necessity modals (e.g. *should*, *ought*) which intuitively express universal quantification over a smaller domain of worlds (von Fintel & Iatridou 2008). See Matthewson 2019 and references therein for further discussion of these two axes.

The examples above show that English modals lexically specify modal force (each modal has a fixed quantificational force) but exhibit variability across flavors (the modals can express more than one flavor). We note that such variability does not require that all modals in English can express all flavors: for instance, *might* arguably can only be used epistemically. Kratzerian semantics for modals capture this by hard-coding quantificational force into the meaning of a modal but relying on context to determine the flavor.[1]

Not all languages are like English: some exhibit so-called *variable force modals*, which specify flavor but do not lexically encode force. This has been found at least in St'át'imcets (Rullmann et al. 2008), Nez Perce (Deal 2011), Old English (Yanovich 2016), and Pintupi-Luritja (Gray 2021).[2] We illustrate the phenomenon with elicited examples of St'át'imcets *k'a* from Rullmann, Matthewson & Davis 2008:[3]

(3) a. [Context: You have a headache that won't go away, so you go to the doctor. All the tests show negative. There is nothing wrong, so it must just be tension.]

nilh	*k'a*	lh(el)-(t)-en-s-wá(7)-(a)	ptinus-em-sút
FOC	INFER	from-DET-1SG.POSS-NOM-IMPF-DET	think-MID-OOC

'It *must* be from my worrying.'

b. [Context: His car isn't there.]

plan	k'a	qwatsáts
already	INFER	leave

'Maybe he's already gone.'

(3a) shows *k'a* being used with strong force and epistemic flavor. (3b) shows *k'a* being used with weak force and epistemic flavor. Further analysis in Rullmann et al. 2008 shows that *k'a* can only be used with epistemic flavor, so it is an example with lexically specified flavor but variable force. The discussed semantic variation across English and St'át'imcets is summarized by Table 1.

1 Typical implementations determine the flavor as the product of two further parameters: a modal base and an ordering source. We set aside this distinction for present purposes.

2 We will discuss modals that specify neither force nor flavor in the next section.

3 These are examples (5c) and and (5e) from Rullmann et al. 2008: 321. See their footnote 5 on p. 320 for the abbreviations.

St'át'imcets *ka*	epistemic	deontic	...
weak		✓	
strong		✓	

English *must*	epistemic	deontic	...
weak			
strong	✓	✓	

Table 1: Two kinds of modal semantic underspecification: *variable-force* and *variable-flavor*.

Even with this substantial cross-linguistic variation, there are kinds of modal meanings that are unattested. We can describe potential restrictions on modal meanings as semantic universals. In order to state universals for modals in a relatively theory-neutral manner (i.e. in a way that does not presuppose a particular formal semantic implementation), we make the following assumptions. We assume that force and flavor are fundamentally properties of contexts of use. This reflects current practice in semantic fieldwork as applied to modality (Matthewson 2004, Bochnak et al. 2020, Vander Klok 2021).[4] For example, the modal questionnaire of Vander Klok 2021 consists exactly of discourse contexts designed to isolate a single force-flavor pair. These contexts can subsequently be used for tasks like elicitation, translation, and acceptability. Finally, we will say that a modal M *can express* a force-flavor pair just in case a bare positive sentence of the form Mp is judged felicitous in a context with that pair.[5]

At this level of generality, we will represent the meaning of a modal as being a set of force-flavor pairs. The semantic universals that we will discuss will be constraints on what kinds of meanings (sets of such pairs) are attested in the languages of the world. For notation, for a modal m, let $[\![m]\!]$ be the set of force-flavor pairs it can express. Furthermore, we will write $\mathrm{fo}(m) = \{\mathrm{fo} \mid \exists \mathrm{fl} \text{ s.t. } (\mathrm{fo}, \mathrm{fl}) \in [\![m]\!]\}$ for the set of forces that a modal m can express and *mutatis mutandis* for $\mathrm{fl}(m)$ and the set of flavors.

We adopt this level of generality because it avoids commitment on the exact formal semantics of these expressions, which is often still being debated. For example, we can say that a *variable force modal* is one that can express more than one pair with the same flavor (i.e. for which $|\mathrm{fo}(m)| > 1$). This is useful because there are two broad approaches to the semantics of such variable force modals: they actually encode existential quantification but lack a universal scalemate (Deal 2011) or they

4 In addition to the particular studies already mentioned, see Matthewson 2013, Cable 2017 for more examples of the application of these methods.

5 We intend 'judged felicitous' to also include the case where such sentences are produced naturally in elicitation tasks, as well as when such sentences are found in naturally-occuring contexts which have a clear force-flavor pair.

encode universal quantification but rely on some mechanism of domain restriction (Rullmann et al. 2008, Bochnak 2015a, Močnik et al. 2019). On such analyses, the underlying semantics contains one specific quantifier; in the present setting, they will still be considered variable force since bare positive sentences are used in contexts with multiple forces.

3 Modal Semantic Universals

Having laid this groundwork, we now introduce the Independence of Force and Flavor (IFF) universal (Steinert-Threlkeld et al. 2023) by showing how it is a refinement of an earlier proposed universal which accommodates recently-discovered counterexamples.

3.1 The Single-Axis of Variability Universal

While the previous section has shown that some modals exhibit variability on the flavor axis (e.g. English *may*) and some modals exhibit variability on the force axis (e.g. St'át'imcets *k'a*), all of the previously discussed expressions are not variable on the other axis. This pattern was observed across many languages from many different families. As a result of a detailed study of the modal systems of six typologically unrelated languages, Nauze 2008 proposed a semantic universal stating that modals cross-linguistically can in fact only exhibit variation along a single axis:

> THE SINGLE AXIS OF VARIABILITY (SAV) UNIVERSAL**:** All modals in natural language satisfy the single axis of variability property: if a modal can express more than one flavor, it can only express one force (and *mutatis mutandis* for force and flavor). That is to say: a modal may exhibit variable force or variable flavor, but not both.[6]
>
> [Alternative formulation: $|\text{fo}(m)| = 1$ or $|\text{fl}(m)| = 1$, where $|\cdot|$ is the set cardinality function.]

At least two counterexamples to this universal have been discovered. The first comes from Washo. Bochnak 2015b,a has argued that the modal verb *-eʔ* can be used in both possibility and necessity contexts with a range of modal flavors. In other words, it exhibits variation both on the force axis as well as the flavor axis. Similarly, Močnik et al. 2019 demonstrate that the Koryak attitude verb *ivək* can be

6 Here is the formulation in Nauze 2008, p. 222: "Modal elements can only have more than one meaning along a unique axis of the semantic space: they either vary on the horizontal axis and thus are polyfunctional in the original sense of expressing different types of modality or they vary on the vertical axis and can express possibility and necessity, but they cannot vary on both axes."

used to express both necessity and possibility. For the doxastic flavor, this means that *ivək* can be used to mean roughly ‘believe’ (necessity) as well as ‘allow for the possibility that’ (possibility). They also argue that the expression can be used to express both doxastic and assertive flavors, thus demonstrating variability on both axes.[7] It is worth noting that while *ivək* exhibits variability along both the force and flavor axes, it is not maximally underspecified: there are still force-flavor combinations that it *cannot* express. Bochnak and Močnik et al. use different variants of the universal quantifier plus choice function analysis of Rullmann et al. 2008 to analyze the respective expressions.

We note also that a refinement of Nauze’s SAV due to Vander Klok 2013b (as reported and discussed in Matthewson 2019) does not accommodate these counterexamples. In particular, Vander Klok proposes that *a modal system as a whole* may only exhibit variability on a single axis in each of the root and epistemic domains. That is: if one root modal exhibits variability on the flavor axis, no other root modal exhibits variability on the force axis (though an epistemic modal may do so) and *mutatis mutandis* for epistemic modals and also for the force axis. This proposal is strictly stronger than Nauze’s: if a language satisfies Vander Klok’s generalization, then every modal therein satisfies SAV. For this reason, counterexamples to the SAV are also counterexamples to this proposal.

3.2 The Independence of Force and Flavor Universal

The counterexamples to the SAV universal show that some languages have modals which are contextually underspecified for *both* force and flavor. It does not follow from this, however, that arbitrary sets of force-flavor pairs are expressed. Intuitively, one does not expect to find a modal in a language that can only express, for instance, epistemic necessity and teleological possibility. Steinert-Threlkeld et al. (2023) use this intuition to define a new semantic universal for modals.

> THE INDEPENDENCE OF FORCE AND FLAVOR (IFF) UNIVERSAL**:** All modals in natural language satisfy the independence of force and flavor property: if a modal can express the pairs $(\mathrm{fo}_1, \mathrm{fl}_1)$ and $(\mathrm{fo}_2, \mathrm{fl}_2)$, then it can also express $(\mathrm{fo}_1, \mathrm{fl}_2)$ and $(\mathrm{fo}_2, \mathrm{fl}_1)$.
>
> [Alternative formulation: a modal m satisfies the IFF property just in case $[\![m]\!] = \mathrm{fo}(m) \times \mathrm{fl}(m)$, where $\times$ is the Cartesian product.]

This universal captures the guiding idea from Kratzer 1981 and much subsequent theorizing on the semantics of modals that force and flavor are *independent*

7 There are also apparently bouletic uses of *ivək*, but Močnik et al. 2019 argue that this flavor does not come from *ivək* alone but from interaction with material in the embedded clause.

axes of meaning. In the standard semantics, this is captured by the separation of quantification from the modal base and ordering source which jointly specify the domain of that quantifier and thereby the flavor. The IFF universal expresses this conception of independence in a theory-neutral way and proposes it as a substantive universal on the semantics of modals cross-linguistically.[8]

One can next ask: why might the IFF universal be true? A slew of recent work has explored domain-general explanations of semantic universals, including accounts from *ease of learning* and *communicative efficiency* (Steinert-Threlkeld et al. 2019, Zaslavsky et al. 2018, Steinert-Threlkeld 2020a, Steinert-Threlkeld et al. 2020, van de Pol et al. 2021, Steinert-Threlkeld 2021, Denić et al. 2022, Uegaki 2021: i.a.). In this work, we explore whether the IFF universal may have emerged due to to the latter; that is, we investigate whether optimally trading off cognitive complexity with communicative accuracy tends to generate languages that more often satisfy IFF. Our aim is thus to illuminate not only whether actually attested modal systems are shaped by efficiency, but also whether there is a systematic relationship between efficiency and 'naturalness'. If efficiency leads to naturalness, this would give a unified account of modal semantic typology: natural language modal meanings result from a general functional constraint to achieve shared goals under bounded cognitive resources.

4 The efficient communication hypothesis

Our working notion of communicative efficiency can be summarized by the following tension. A language can be simple and uninformative (e.g. containing a single expression). A language can be complex and informative (e.g. containing unique expressions for each possible thought to be expressed). A language cannot be both simple and informative: these two pressures trade-off against each other. A hypothesis in linguistics is that the natural languages are (near) solutions to this multi-objective optimization problem, and that these efficiency pressures explain constraints on crosslinguistic variation (Kemp et al. 2018).

This efficient communication hypothesis has been successfully applied across a variety of semantic domains including kinship terms, color terms, number terms, container terms, quantifiers, boolean connectives, indefinite pronouns and deictic adverbs, among others (Kemp & Regier 2012, Regier, Kemp & Kay 2015, Xu & Regier 2014, 2016, Steinert-Threlkeld 2021, Uegaki 2021, Denić, Steinert-Threlkeld & Szymanik 2022, Chen, Futrell & Mahowald 2022). We follow others in this literature in using a computational experiment to simulate the simplicity/informativeness

8 One sense in which the formulation can be seen as 'theory-netural': Kratzer 1981 builds in independence by treating force as lexically encoded and flavor as contextually determined. The present level of analysis does not commit to any positive view on which components are lexically specified and which are not. Thanks to Wataru Uegaki (p.c.) for discussion here.

trade-off. Generally, if the natural languages are optimal solutions to the trade-off (or closer to being optimal than non-natural languages), this suggests that pressures for communicative efficiency cause the observed patterns for that semantic domain.

To show that the natural language modal inventories and the semantic universals that hold of them are indeed shaped by such general pressures, we require measures of efficiency that are appropriate to modal semantics. In the remainder of this section, we describe our measures of simplicity and informativeness in detail. Henceforth, we will use the term 'language' to mean a modal inventory, i.e. a set of modal meanings.[9]

4.1 Simplicity

We define simplicity in terms of its inverse, complexity. We model the complexity of a modal meaning as the fewest number of atoms it takes to express its meaning in a Language of Thought (LoT) (Fodor 1975, Feldman 2000, Goodman et al. 2008, Piantadosi et al. 2016, Denić et al. 2022). This representation language is the standard language of propositional logic (see Appendix A.1). The language has an atom both for each flavor and for each force. The primitive operators in the language include conjunction ($\wedge$), disjunction ($\vee$), and negation ($\neg$) of features. As an example, in this language, we can express the meaning of English *might*, $[\![might]\!] = \{(\text{weak}, \text{epistemic})\}$ as $w \wedge e$, where w is the atom for weak force and e is the atom for epistemic flavor.

We extend heuristics described in Feldman 2000 to find the shortest boolean formula for a modal in this language (see Appendix A.2). This allows us to map any modal meaning to a discrete measure of its complexity, using a collection of the points it can express. In particular, we write down a disjunctive normal form (DNF) expressing the disjunction of all pairs that a modal can express, and then apply an algorithm to shorten this DNF. A key rule in the minimization algorithm applies the fact that conjunction distributes over disjunction, allowing one to replace a formula like $(w \wedge e) \vee (w \wedge d)$ with a formula like $w \wedge (e \vee d)$.[10] This is intended to capture the intuition that some meanings differ in terms of in how difficult it is to compactly represent their variability on the two axes. In particular, when features of meaning share an axis, this axis may be 'factored' out in the shortest formula representation. Some example applications of the minimization algorithm are illustrated in Table 2. The pseudocode for this algorithm and the rules of inferences it uses are described in Appendix A.

9 More precisely, a language is a multi-set of modals, which allows for synonymy.

10 The Quine-McCluskey algorithm is a standard minimization algorithm (Quine 1952), but it only produces minimal disjunctive normal forms. The rules that we apply can produce shorter formulas that are not in such form.

Modal	Meaning representation					Shortest Formula in LOT	Complexity (# of atoms)
		e	*d*	*c*	*t*		
may	*w*	✓	✓			$w \wedge (e \vee d)$	3
	s						
		e	*d*	*c*	*t*		
mought	*w*	✓				$(w \wedge e) \vee (s \wedge d)$	4
	s		✓				
		e	*d*	*c*	*t*		
notcirc	*w*	✓	✓		✓	$\bar{c}$	1
	s	✓	✓		✓		

Table 2: Measuring complexity for English *may* and two hypothetical modals *mought* and *notcirc*. First column: meaning representation. Second column: shortest LOT formula. Third column: complexity measure.

Given this measure of the complexity of any modal in isolation, we can measure the overall complexity of a language as a sum of the complexities of the modals therein. Formally:

$$\mathrm{Comp}(L) := \sum_{m \in L} \min\{\mathrm{len}(\varphi) : \varphi \in \mathrm{LOT}, [\varphi] = [\![m]\!]\}$$

For example, if a language consisted of exactly one of each of the modals in Table 2, it would be assigned $\mathrm{Comp}(\textit{may}) + \mathrm{Comp}(\textit{mought}) + \mathrm{Comp}(\textit{notcirc}) = 8$. To summarize, we have used a minimum description length approach to quantify the complexity of languages as a sum of the complexities of each of the items in its modal vocabulary.

4.2 (Literal) Informativeness

The informativeness of a language is modeled after the idea of successful communication of signals between literal speakers and listeners (Skyrms 2010, Steinert-Threlkeld 2021). This measure can be modeled as an expected utility of a language L for communication, where the expectation is taken over repeated interactions between a speaker trying to successfully convey a force-flavor pair $p \in P$ to a listener. More precisely:

(1)
$$\begin{aligned} I(L) &:= \mathbb{E}[u(p, p')] \\ &= \sum_{p \in M} \mathbb{P}(p) \sum_{m \in L} \mathbb{P}(m|p) \sum_{p' \in m} \mathbb{P}(p'|m) \cdot u(p', p) \end{aligned}$$

In Equation 1, $\mathbb{P}(m|p)$ is the probability a speaker selects a specific modal m to communicate a meaning p (a single (fo, fl) pair in the semantic space). $\mathbb{P}(p'|m)$ is the probability that a listener guesses a (fo, fl) pair p', given the expression heard (m). A prior over meaning points $\mathbb{P}(p)$ models how often agents need to communicate about specific meanings. In principle, different linguistic communities will have different communicative need distributions. We estimate one distribution from English corpus data (Pyatkin et al. 2021) to measure informativeness for every language in our main experiment (details in Section 5.2), leaving estimation of the communicative need distributions of the other languages to future work.

The utility function $u(p, p')$ measures how 'good' the listener's guess p' is, if the speaker intended to convey p. The structure present in the modal meaning space allows us to communicative utility as a graded notion, with some utility awarded to guesses that are better than others. In particular, we define a utility scoring function $u(p, p')$ which gives half-credit (0.5) to correctly guessing each of the force and the flavor of p. Thus, if p' shares one axis of meaning with p, the utility will be 0.5; if it shares both, 1; and if it shares neither, 0. More precisely:

(2) $$u(p, p') = 0.5 \cdot \mathbf{1}\{\text{fo}(p) = \text{fo}(p')\} + 0.5 \cdot \mathbf{1}\{\text{fl}(p) = \text{fl}(p')\}$$

where $\mathbf{1}\{x\}$ is the indicator function which returns 1 if x is true, and 0 if x is false.

Lastly, just as we measure complexity instead of simplicity, we define the *communicative cost* of a language as the 'inverse' of its informativeness: $Cost(L) = 1 - I(L)$. In other words, while simplicity and informativeness are "desirable" features for a language, complexity and communicative costs are "undesirable" features: they should both be *minimized* to the extent possible.

5 Computational experiment

In order to evaluate the simplicity/informativeness trade-off for modals, we will measure the *optimality* of a language as the distance to the optimal solutions along an estimated Pareto frontier. The Pareto frontier is a set of optimal solutions to a trade-off problem, which in this setting corresponds to the set of languages which achieve the minimum complexity for a given value of communicative cost. We will also measure the degree to which each language satisfies IFF to see if this variable—which we will call *naturalness*—correlates with optimality. Accordingly, the experiment involves the following steps: we (1) collect a sample of natural language modal inventories to measure; (2) fix a semantic feature space from which to generate meanings; (3) find the shortest expression for each meaning; (4) estimate the communicative need distribution over meanings; (5) estimate the Pareto frontier; (6) generate a sample of hypothetical modal systems; and (7) measure the optimality of natural and hypothetical languages by each language's distance to the frontier. We describe these steps in turn in the next section, before presenting our main results. The code for reproducing these results can be found at https://github.com/nathimel/modals-effcomm .

5.1 Typological Data

To measure the communicative efficiency of natural language modal inventories, we use the recently introduced Database of Modal Typology (Guo et al. 2022). This is a public repository for linguists to contribute data they have collected on crosslinguistic modal semantics. Individual modal expressions are annotated for force and flavor, among other linguistic features. At time of writing, the database contains 40 languages; of these, we consider only the languages for which linguists indicated the modal inventories were described completely, with positive and negative truth value judgments from speakers in elicitation tasks. This restriction results in a total of 27 languages from 17 families that we use for our analysis, described in Table 3. This filtered sample also constrains the maximum language size and meaning space used in the experiment (more details in Section 5.2). A sample of the Tlingit data recorded in the database and used in our analyis is given in Table 4.

All the languages we measure in this work are represented as a collection of their modal expressions, and in turn each expression represented as a collection of the force-flavor pairs it can be used to communicate. The natural languages are obtained from the database, while the hypothetical ones are mathematically generated. We describe this generation procedure in the next subsection.

5.2 Methods

Meaning space Our main results center a meaning space with two modal forces $\{\text{weak}, \text{strong}\}$ and three modal flavors $\{\text{epistemic}, \text{deontic}, \text{circumstantial}\}$, for a total of six possible meaning points. This meaning space was chosen because it is the maximal set of force-flavor pairs that could be expressed by the natural language modal inventories (i.e., the intersection of the languages' observed meanings). The possibility of modeling languages with different 'domains' of modality (possibly hierarchically structured with, e.g. root/epistemic as being fundamental) will be left for future work.

Shortest expressions There are $2^6 - 1 = 63$ modal meanings in this space (non-empty sets of force/flavor pairs). For each of these, we apply the minimization algorithm described in Section 4.1 to find the shortest formula expressing that meaning, thereby determing the complexity of each modal meaning.

Name	Family	Source
Akan	Atlantic-Congo	Uegaki et al. (2022)
Basque	Basque	Uegaki et al. (2022)
Cantonese	Sino-Tibetan	Uegaki et al. (2022)
Central Khmer	Austroasiatic	Uegaki et al. (2022)
Dutch	Indo-European	Uegaki et al. (2022)
Gitksan	Tsimshian	Matthewson (2013)
Goemai	Afro-Asiatic	Hellwig (2011)
Hausa	Afro-Asiatic	Uegaki et al. (2022)
Modern Hebrew	Afro-Asiatic	Uegaki et al. (2022)
Hindi	Indo-European	Uegaki et al. (2022)
Hungarian	Uralic	Uegaki et al. (2022)
Javanese	Austronesian	Vander Klok (2013a)
Korean	Koreanic	Uegaki et al. (2022)
Lillooet	Salishan	Rullmann et al. (2008)
Logoori	Atlantic-Congo	Gluckman et al. (2020)
Mandarin	Sino-Tibetan	Uegaki et al. (2022)
Modern Greek	Indo-European	Uegaki et al. (2022)
Tundra Nenets	Uralic	Nikolaeva (2014)
Turkish	Turkic	Uegaki et al. (2022)
Igbo	Atlantic-Congo	Uegaki et al. (2022)
Japanese	Japonic	Uegaki et al. (2022)
Russian	Indo-European	Uegaki et al. (2022)
Tagalog	Austronesian	Uegaki et al. (2022)
Tharaka	Atlantic-Congo	Uegaki et al. (2022)
Tlingit	Athabaskan-Eyak-Tlingit	Cable (2017)
Vietnamese	Austroasiatic	Uegaki et al. (2022)
Western Farsi	Indo-European	Uegaki et al. (2022)

Table 3: The 27 languages used in our computational experiment, taken from the modal typological database described in Guo et al. 2022. The data for these languages result from published and unplished elicitation tasks and span 17 distinct families.

expression	force	flavor	can_express
giwe	weak	circumstantial	0
shákdé	weak	circumstantial	0
future mode	weak	circumstantial	0
potential mode	weak	circumstantial	1

Table 4: Example of our basic data format for several strategies of expressing modality in Tlingit (Cable 2017).

Communicative need distribution We estimate one communicative need distribution over force-flavor pairs to measure the informativeness of all languages in the experiment. This is done using a fine-grained annotation of the Georgetown Gradable Modal Expressions (GME) Corpus, as described in Pyatkin et al. (2021). The GME is an expert-annotated sample of documents included in the Opinion Corpus introduced by Wiebe et al. (2005), which consists of roughly ten thousand sentences drawn from articles in English from the world press published during 2001 and 2002. We use the "Fine-Grained" annotation category of data from Pyatkin et al. (2021) which uses a taxonomy of modal flavors compatible with our chosen feature space. The communicative need distribution is the relative frequencies of each force-flavor pair observed in the corpus. To obtain these frequencies, we count the occurrences of modal verbal auxiliaries. The resulting distribution is displayed in Table 5.

	epistemic	deontic	circumstantial
weak	0.139	0.042	0.143
strong	0.104	0.254	0.318

Table 5: Estimated probability distribution over force-flavor pairs, representing a communicative need distribution over modal meanings.

Estimating the Pareto frontier To estimate the Pareto frontier of languages that optimally balance complexity and communicative cost, we apply an evolutionary algorithm to directly optimize these two objectives (Steinert-Threlkeld 2020b, 2021, Denić et al. 2022). This works as follows. In the beginning, a seed population of 2000 artificial modal languages is randomly generated (using the first sampling procedure described in the next section). There are then several (200) 'generations'. At the end of each generation, a random choice of between 1 and 5 mutations is applied to each of the dominant languages. These dominant languages represent the subset of their generation best optimizing the simplicity/informativeness trade-off. The mutations include randomly adding a modal to a language, removing a modal from a language, and replacing a modal in a language. Another mutation removes a single force/flavor pair from the meaning of a given modal in a language, and the last mutation adds to a language one modal that can express only a single force/flavor pair. Each dominant language has enough 'offspring' via mutation to create 2000 languages at each generation. After 200 generations of this process, the dominant languages are the estimated Pareto frontier. [11]

Sampling languages In addition to measuring the modal inventories of natural languages, we also generate a large sample of hypothetical modal inventories. Because exhaustive enumeration of the space of possible modal languages is not feasible,[12] we use several sampling techniques to encourage exploration of the space of possible languages.

Our sampling procedure has two steps. First, one sample of languages is obtained from random/unbiased sampling. We manipulate both the size of the language (from one to ten modals) and the number of modals in the language satisfying the IFF universal (from one to the current size). The maximum vocabulary size was chosen for feasibility reasons, but also because the largest vocabulary in our typological data contains ten modals. For each combination of these two parameters, we generate languages by sampling mutisets of modal expressions from the set of possible expressions that do and do not satisfy IFF. Because of limitations on how many unique languages exist for each combination, attempting to generate 40000 languages with equal representation of each combination results in 31524 total languages.

Second, to encourage significant exploration of the space of possible languages, especially the low-density regions unlikely to be discovered by the above random sampling procedure, we apply the same evolutionary algorithm for estimating the Pareto frontier of efficient languages three more times: once for each of the other corners of the two-dimensional (complexity, communicative cost) space of possible languages. In other words, while the main evolutionary algorithm sought to *minimize*

11 The full details of this algorithm, including pseudocode, can be found in Appendix A of Steinert-Threlkeld (2021). In addition to the three mutations applied to entire languages in that algorithm (namely, `add`, `remove`, and `interchange`), we use two mutations that apply at the level of individual lexical items: `add_point` and `remove_point`. The `add_point` mutation adds a random force-flavor pair to the meaning of a randomly chosen modal expression within a language. Similarly, the `remove_point` mutation removes a force-flavor pair from a modal.

12 In our meaning space with two forces and three flavors, there are six meaning points, yielding $2^6 - 1 = 63$ possible modals. For a fixed vocabulary size of 10 modals (allowing for synonymy), there are $\binom{k+n-1}{k} = \binom{10+63-1}{10} > 5.3 \times 10^{11}$ possible languages.

both measures, in order to encourage exploration, we look at all combinations of minimizing and maximizing both measures.

We combine all the unique languages discovered in the experiment, (i) by random sampling and (ii) each of the four runs of the evolutionary algorithm, into one pool of languages. Including the 27 natural languages, we obtain a sample of $75,000$ total languages for our analysis.[13]

Measuring optimality Finally, each language is measured for complexity, communicative cost, and *optimality*. We define the optimality of a language to be its minimum Euclidean distance to a point on the Pareto frontier.

Hypotheses We measure *naturalness* as a continuous value of languages (the fraction of the vocabulary conforming to the IFF universal). If the modal inventories in natural language have been shaped by the simplicity/informativeness trade-off, then we expect that the natural languages will be among the optimal languages, and that optimality will be significantly correlated with the IFF degree.

5.3 Results

Figure 1 depicts the main results. This plot shows all of the $N = 75000$ languages generated as described above on a two-dimensional plane, with the x-axis being complexity and the y-axis being communicative cost. The black circles are the languages discovered during our search procedure that balance complexity and communicative cost better than any other languages. These languages constitute the estimated Pareto frontier. Each natural language is labeled and marked by a red '+'. The size of the dots correspond to the number of distinct languages mapped to that point in this complexity/cost space. The color corresponds to the naturalness, i.e. what percentage of the modals in the language satisfy IFF (with 0.0 being blue and 1.0 being yellow).

We catalogue three main results. First, all Pareto-optimal languages are 'perfectly natural'. That is, every estimated solution to the simplicity/informativeness trade-off is a lexicon containing only modals that satisfy the IFF universal.

This leads to our second main observation: naturalness is highly correlated with simplicity, informativeness, and Pareto-optimality. The strength of some of these relationships can be seen observed from visual inspection: as languages become 'more natural' / 'more IFF', they also become simpler, more informative, and

13 We chose $75,000$ as the size of the sample languages used for final analysis to support more intuitive comparisons across experiment parameters (described in Section 6), as such parameter variations cause our sampling procedure to encounter different numbers of languages.

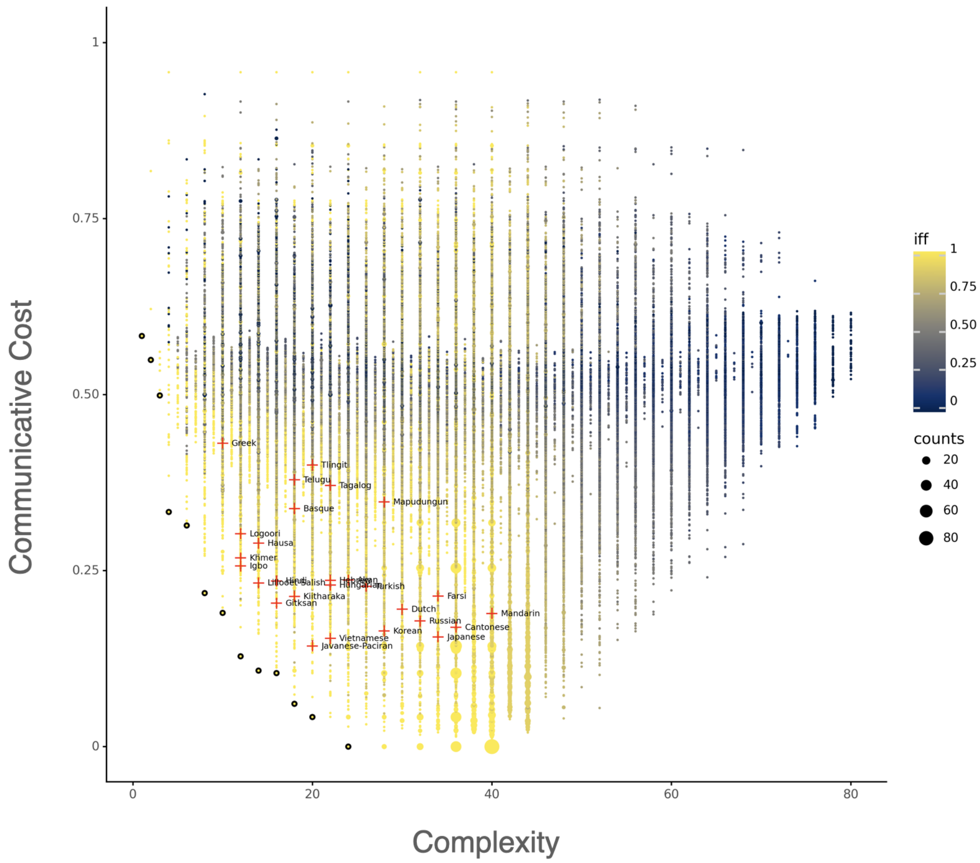

Figure 1: The complexity/communicative cost trade-off for modal languages. The black circles constitute the estimated Pareto frontier of optimal solutions to the trade-off. Size corrresponds to the number of languages at a given point in the trade-off space. Red '+' are natural languages obtained from Guo et al. 2022 and are labeled. The color of a language is its degree of naturalness (proportion of modals satisfying the IFF universal).

language	N	simplicity	informativeness	optimality
mean natural	27	.722	.750	.894
mean population	75000	.545	.556	.731

Table 6: Mean simplicity, informativeness, and Pareto optimality for attested natural languages vs hypothetical languages.

closer to the Pareto frontier. Measuring Pearson correlations shows that a language's degree of naturalness is strongly correlated with both simplicity ($r(75k) = .35$) and informativeness ($r(75k) = .44$), and most of all with optimality ($r(75k) = .48$). This suggests that typological constraints on modal systems may emerge from optimally *trading off* complexity with communicative cost, rather than from optimizing either alone.

Finally, our third observation regards the efficiency of actual modal semantic systems. The 27 attested systems are closer on average to the Pareto frontier than merely mathematically possible, artificially generated systems. Even the languages that are further from the frontier (e.g., Cantonese, Mandarin, Japanese, Farsi, Mapudungun) appear to be significantly more efficient than many hypothetically possible languages. We return to a discussion of these less efficient attested systems in Section 7. A comparison of mean simplicity, informativeness and Pareto optimality for the natural languages and the generated hypothetical languages is given in Table 6. This table shows that the natural languages are simpler, more informative, and also more optimal than merely possible languages.[14]

Before moving forward to a general discussion of these results, we first ask: how well preserved are these trends under different ways of measuring efficiency? In the next subsection, we address this question by exploring the robustness of the main results to alternative measures of informativeness.

6 Robustness to alternative measures of informativeness

The informativeness measure has three key components: the communicative need distribution, the speaker/listener distributions, and the communicative utility function. In what follows, we define alternative choices for each of these components, and then show that the main results just reported are not significantly affected by these choices.

14 Given the large sample sizes in our data, we do not report significance tests for these means or the aforementioned correlations, which are sensitive thereto.

6.1 Communicative need: estimated vs. uniform

Our main result was obtained using a prior over force-flavor meaning pairs that was estimated from a corpus. Since this estimate comes from one corpus within a single genre, it is unclear how representative it is as an accurate model of communicative need. To further explore the robustness of our results, we will also repeat our analysis using a uniform distribution.

6.2 Communicative utility: graded vs. binary

While our first experiment reported results using a graded notion of utility for modal-meanings (Equation 2), there are other ways of modeling a successful communicative interaction. Another reasonable criterion might require that only perfect transmission of a force-flavor pair should be counted as success. To this end, we define an additional utility function $u_{\text{bin}}(p, p')$ which awards success for a communicative interaction if and only if $p = p'$. We then have u_{bin} simply be an indicator function:

$$u_{\text{bin}}(p, p') = \mathbf{1}\{p = p'\}.$$

6.3 Literal vs. pragmatic speakers

So far, we have described how to measure the communicative utility of modal lexicons with respect to *literal* speakers and listeners. We now describe a measure of informative communication relative to *pragmatic* speakers and listeners, using models from the Rational Speech Act (RSA) framework (Frank et al. 2012, Degen 2023). The basic intuition for the RSA model is as follows. While a literal listener does not weigh the benefits or costs of using particular expressions to communicate their intended meaning, a pragmatic speaker considers the utility of different expressions depending on the way a literal listener will interpret their utterance.[15] A pragmatic listener chooses to interpret a modal expression based on how a pragmatic speaker would choose to utter the expression. These models are formalized in Equations 3 and 4:

15 Here we describe the S_1 speaker and L_1 listener, which index the ‘first level’ of recursive reasoning, but in general the RSA framework can consider pragmatic reasoning of agents up to arbitrary recursion depth. For example, an S_2 speaker can condition their decisions based on how a pragmatic listener would interpret them, and so on.

$$(3) \qquad \mathbb{P}_{\text{prag}}(m|p) = \frac{\exp(\alpha \cdot U_{\text{prag}}(m,p))}{\sum_{m' \in M} \exp(\alpha \cdot U_{\text{prag}}(m',p))}$$

$$(4) \qquad \mathbb{P}_{\text{prag}}(p|m) = \frac{\mathbb{P}(m|p) \cdot \mathbb{P}(p)}{\sum_{p' \in M} \mathbb{P}(m|p')},$$

where the utility for a pragmatic speaker to choose a modal expression m to communicate a force-flavor pair p is defined in Equation 5. This utility is the log probability of the literal speaker's interpreting modal m as meaning p:

$$(5) \qquad U_{\text{prag}}(m,p) = \log \mathbb{P}(m|p).$$

The idea behind this pragmatic assignment of utility is that the more likely that an expression will be interpreted correctly by a listener, the better it is for a speaker to use it.[16] The pragmatic speaker chooses modal expressions in proportion to their utility.[17] Meanwhile, the pragmatic listener selects the Bayesian-optimal force-flavor interpretation, given their knowledge about what modals mean in their language, and a prior over force-flavor pairs. With these ingredients in place, we now define the *pragmatic informativeness* of a language L in Equation 6:

$$(6) \qquad I_{\text{prag}}(L) = \sum_{p \in M} \mathbb{P}(p) \sum_{m \in L} \mathbb{P}_{\text{prag}}(m|p) \sum_{p' \in M} \mathbb{P}_{\text{prag}}(p'|m) \cdot u(p,p'),$$

where u can be either graded or binary communicative utility. Mirroring our definition of literal communicative cost (Equation 1), we define pragmatic communicative cost as $Cost_{\text{prag}}(L) = 1 - I_{\text{prag}}(L)$.

6.4 Results

For each combination of these three different manipulations—literal vs. pragmatic speakers and listeners, binary vs. graded communicative payoff, and an estimated vs.

16 For example, if the literal listener never interprets English *may* as expressing $(\text{strong}, \text{circumstantial})$, then $\mathbb{P}(\text{may}|(\text{strong}, \text{circumstantial})) = 0$ and $U_{\text{prag}}(\text{may}, (\text{strong}, \text{circumstantial})) = -\infty$, i.e. the pragmatic speaker should never use *may* to express this force-flavor pair. Meanwhile, if the literal listener always interprets English *might* as expressing $(\text{weak}, \text{epistemic})$, then $\mathbb{P}(\text{might}|(\text{weak}, \text{epistemic})) = 1.0$ and so $U_{\text{prag}}(\text{might}, (\text{weak}, \text{epistemic})) = 0.0$, i.e., saying *might* will maximize the pragmatic speaker's utility.

17 This is modeled by the softmax function in Equation 3, which normalizes the utility score of each modal expression to the probability of uttering the expression. The temperature parameter α in this function controls the 'sharpness' of the resulting distribution. When $\alpha \to \infty$, the speaker's choice of expression depends only on its utility; when $\alpha = 0$, the speaker chooses randomly (and hence will typically be much worse for communication than a literal speaker). In our experiment, we let $\alpha = 1.0$, which reflects the assumption that speakers are neither perfectly rational nor completely incompetent.

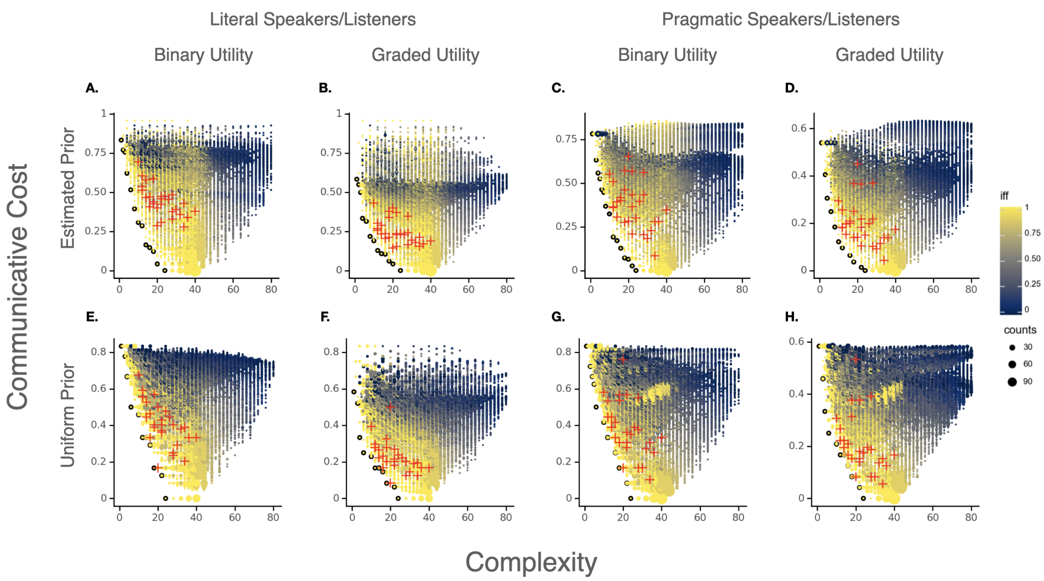

Figure 2: The complexity/communicative cost trade-off for modal languages, evaluated with eight parameter variations for measuring communicative cost. Note that subplot (**B**) is a reproduction of Figure 1. Subplots (**A** - **D**) are obtained with the estimated prior over meanings, while (**E-H**) are obtained assuming a uniform prior. (**A/E**) The results obtained for literal speakers and listeners, with respect to a binary utility score for communicative interactions. (**B/F**) Results assuming literal speakers and listeners and a graded communicative utility function. (**C/G**) Results assuming pragmatic speakers and listeners and binary utility. (**D/H**) Results obtained for pragmatic speakers and listeners and graded utility.

Prior	Speaker/Listener	Utility	Property	IFF corr.	$\overline{\text{natural}}$	$\overline{\text{population}}$	$\overline{\text{nat.}} - \overline{\text{pop.}}$
estimated	literal	binary	simplicity	.460	.722	.515	.207
			informativeness	.506	.560	.391	.169
			optimality	.566	.877	.709	.168
		graded	simplicity	.348	.722	.545	.177
			informativeness	**.437**	**.750**	**.556**	**.194**
			optimality	**.479**	**.894**	**.731**	**.163**
	pragmatic	binary	simplicity	.578	.722	.458	.264
			informativeness	.307	.630	.427	.203
			optimality	.625	.886	.667	.219
		graded	simplicity	.667	.722	.365	.357
			informativeness	.569	.800	.604	.196
			optimality	.765	.889	.587	**.302**
uniform	literal	binary	simplicity	.401	.722	.564	.168
			informativeness	.675	.590	.395	.195
			optimality	.599	.906	.752	.154
		graded	simplicity	.421	.722	.539	.183
			informativeness	**.694**	**.772**	**.567**	.205
			optimality	**.617**	**.914**	**.729**	.185
	pragmatic	binary	simplicity	.372	.722	.576	.146
			informativeness	.510	.618	.453	.165
			optimality	.552	.907	.770	.137
		graded	simplicity	.570	.722	.445	.277
			informativeness	.606	.790	.594	.196
			optimality	.734	.908	.629	**.279**

Table 7: Quantitative evaluation of the efficiency of modal semantic systems for different models of communicative informativeness. Each analysis (separated by a light gray line) measures a total sample of 75000 languages. Column 1 (**Prior**): type of assumed communicative need distribution (prior over meanings). Column 2 (**Speaker/Listener**): type of assumed speaker/listener. Column 3 (**Utility**): type of assumed communicative utility function. Column 4 (**Property**): The property measured of a language. Column 5 (**IFF corr.**): Pearson *r* correlation coefficient of the property with naturalness (degree-IFF). Column 6 ($\overline{\textbf{natural}}$): mean value of the property with respect to the 27 attested languages. Column 7 ($\overline{\textbf{population}}$): mean value of the property with respect to the total sample of 75000 languages. Column 8 ($\overline{\textbf{nat.}} - \overline{\textbf{pop.}}$): the difference betweeen a property's mean value in the attested systems and in the total sample. Bolded values indicate notable contrasts that are discussed in the main text.

uniform prior over meanings—we repeated our main experiment, resulting in eight total analyses. The predicted complexity/communicative cost trade-offs are depicted in Figure 2. For all variations on informativeness, each of the three trends from the main results are reproduced: (1) every Pareto-optimal system is perfectly natural (has degree-1.0 IFF satisfaction); (2) a language's degree of naturalness correlates with simplicity, informativeness, and Pareto-optimality; and (3) the attested modal systems are more Pareto-optimal than the vast majority of hypothetical systems.

We report quantitative evaluations of modal vocabulary efficiency in Table 7.[18] We catalogue several additional takeaways. First, assuming a uniform prior and literal communication, naturalness is more correlated with informativeness than with optimality. Second, however, assuming a uniform prior over meanings, the modal typology generally appears to be even more strongly shaped for efficiency than assuming the prior estimated from Pyatkin et al. (2021). Some of these trends can be observed visually in Figure 2. For example, in the configuration most similar to the one we initially introduced (graded communicative utility, literal speakers and listeners), attested natural languages are higher in informativeness and optimality, and degree-naturalness is more strongly correlated with these measures. This contrast is bolded in Table 7 for comparison. Also bolded are the greatest differences in optimality between the attested and total population of languages, under an informativeness model of graded utility and pragmatic agents. These values suggest that, comparing attested languages to hypothetical languages, the relative gain in optimality is maximized under the more 'intuitive' notion of communication: that is, a pragmatic task with varying degrees of success.

Overall, while there is subtle variability our modeling predictions, the major trends and general picture are stable: the modal typology optimizes for communicative efficiency, and a shared property of all natural language modals—the IFF semantic universal—appears to emerge from this optimization.

7 Discussion

Let us take stock. We set out to address whether modal typology can be accounted for in terms of pressures for efficient communication. To do this, we defined a complexity measure for modal semantic systems, motivated from the idea that the complexity of a mental represenation will be correlated with its shortest description in a Language of Thought. We also considered eight different measures of informativeness, varying (i) our assumptions about the communicative need distribution over meanings, (ii) the notion of what counts as a useful communicative interaction, and (iii) what kind

18 Note that although we do not explicitly manipulate complexity, by manipulating the informativeness objective, our sampling procedure (described in Section 5.2) encounters different languages, which results in different distributions of complexity across languages in each sample.

of reasoning speakers and listeners of a language engage in during communication. We collected 27 attested modal inventories from 17 different language families, and situated facts about their typological variation in terms of a recently proposed lexical semantic universal, the Independence of Force and Flavor (IFF). We then performed computational experiments to evaluate whether typological variation in the world's attested modal semantic systems, and the modal semantic universal, can be explained by pressures for communicative efficiency.

We found that communicative efficiency, operationalized in terms of a joint optimization of competing pressures for simplicity and informativeness, explains much of the observed variation in modal semantic typology. Specifically, we found that (1) every Pareto-optimal solution to the simplicity/informativeness trade-off perfectly satisfies the IFF universal; (2) the degree of naturalness for a language, measured in terms of the number of IFF modals it has, is highly correlated with optimality; and (3) the attested modal inventories are more optimal than the vast number of hypothetically possible modal inventories.

Pressure for efficient communication, in the simplicity/informativeness trade-off sense, appears to explain important variation in modal typology. It surely does not explain all of it. Importantly, none of the modal systems in our sample were found to perfectly optimal, and many have substantial distance from the Pareto frontier. One reason that these systems are further from the frontier is that they may have high degrees of *synonymy*: languages contain multiple modals that encode the same meaning. Mandarin is one such language, containing ten modals overall: three modals that all 'mean' (i.e. have as their minimal LoT formula) $\texttt{strong} \wedge \texttt{epistemic}$, three modals for $\texttt{strong} \wedge \neg\texttt{epistemic}$, two modals for $\texttt{weak} \wedge \texttt{epistemic}$, and two modals for $\texttt{weak} \wedge \neg\texttt{epistemic}$. Such synonymy always hurts the trade-off: it adds complexity without conferring any additional usefulness in informativity.

This raises a number of interesting possibilities. We suspect that modals that we treat as absolute synonyms are lexical items that differ along additional dimensions of meaning that are simply not captured by our coarse-grained semantic features. Recall that we only consider weak and strong modal force, and three modal flavors. Furthermore, while we assume every language has *some* means of expressing different modal notions, it is not the case that they will have designated, independent lexical items to do so. They may rely on more complex compositional constructions. For example, Urdu/Hindi expresses modality via a specific set of morpho-syntactic constructions (Bhatt et al. 2011). Additionally, Tlingit has few grammatical strategies to express modality, and instead uses various pragmatic strategies to communicate about its modal categories (Cable 2017).

It is important to acknowledge that there are many factors shaping semantic typology, efficient communication just being one particularly general and powerful one. It will be important in future work to develop methods to adjudicate between alternative explanations, as well as to distinguish between conceptual/cognitive forces shaping semantic variation from cultural/historical/sociological forces. We have focused on the former because semantic universals seem especially likely to arise from such general cognitive pressures.

There are several salient directions for future research. As previously mentioned, a more detailed efficient communication analysis should (i) represent finer-grained distinctions in meaning than the high-level force-flavor dichotomy that we have considered here, (ii) account for the interactions of modality with other domains of meaning (e.g., tense, aspect, evidentiality), and (iii) measure the contribution of morphosyntactic complexity to the overall complexity of modal expressions (Mollica et al. 2021, Carcassi et al. 2023, Denić et al. 2023). In addition, while we have argued (as have many others) that a part of the lexicon is optimized for efficient communication, it remains less clear how languages optimize the simplicity/informativeness trade-off in cultural evolutionary time. A diachronic analyis of the efficiency of modals (as for example in Zaslavsky et al. (2022)) will be crucial to evaluate whether natural language modal systems actively evolve under pressure for efficiency.

Altogether, our empirical results suggest that efficient communication successfully explains significant variation in modal typology. More generally, our findings lend support to the idea that the humanly-possible languages are the outcomes of independent constraints that arise from representing the environment and coordinating on shared goals. In addition, by explaining how greater optimality results in higher degrees of naturalness, our account is compatible with a view that semantic universals correspond to soft rather than hard constraints. This leaves room for the attractive idea that linguistic 'universals' can be observed as robust, statistical tendencies, rather than absolute restrictions (Evans et al. 2009). Here, we have used an efficiency-based analysis to explain why languages may satisfy these tendencies to varying degrees: modal vocabularies are subject to varying, independent pressures for cognitive simplicity and informative communication.

References

Bach, Emmon & Wynn Chao. 2009. 8 On Semantic Universals and Typology. In Morten H. Christiansen, Christopher Collins & Shimon Edelman (eds.), *Language Universals*, 0. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195305432.003.0008. https://doi.org/10.1093/acprof:oso/9780195305432.003.0008 (19 June, 2023).

Barwise, Jon & Robin Cooper. 1981. Generalized Quantifiers and Natural Language. *Linguistics and Philosophy* 4(2). 159–219. https://doi.org/10.1007/BF00350139.

Bhatt, Rajesh, Tina Bögel, Miriam Butt, Annette Hautli & Sebastian Sulger. 2011. Urdu/Hindi modals. *LFG '11 Conference*. 47–67. http://web.stanford.edu/group/cslipublications/cslipublications/LFG/16/papers/lfg11bhattetal.pdf.

Bochnak, M Ryan. 2015a. Underspecified modality in Washo. In Natalie Weber & Sihwei Chen (eds.), *Workshop on Structure and Constituency in Languages of the Americas 18 & 19*, vol. 39 (University of British Columbia Working Papers in Linguistics), 3–17.

Bochnak, M Ryan. 2015b. Variable force modality in Washo. In Thuy Bui & Deniz Özyıldız (eds.), *North-East Linguistic Society (NELS) 45*, 105–114.

Bochnak, M. Ryan & Lisa Matthewson. 2020. Techniques in Complex Semantic Fieldwork. *Annual Review of Linguistics* 6(1). 261–283. https://doi.org/10.1146/annurev-linguistics-011619-030452.

Cable, Seth. 2017. The Expression of Modality in Tlingit: A Paucity of Grammatical Devices. *International Journal of American Linguistics* 83(4). 619–678. https://doi.org/10.1086/692975.

Carcassi, Fausto & Giorgio Sbardolini. 2023. Assertion, denial, and the evolution of Boolean operators. *Mind & Language*. https://doi.org/10.1111/mila.12448. (30 June, 2023).

Chen, Sihan, Richard Futrell & Kyle Mahowald. 2022. Investigating information-theoretic properties of the typology of spatial demonstratives. In *Proceedings of the 4th workshop on research in computational linguistic typology and multilingual NLP*, 94–95. Seattle, Washington: Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.sigtyp-1.12. https://aclanthology.org/2022.sigtyp-1.12.

Deal, Amy Rose. 2011. Modals Without Scales. *Language* 87(3). 559–585. https://doi.org/10.1353/lan.2011.0060.

Degen, Judith. 2023. The Rational Speech Act Framework. *Annual Review of Linguistics* 9(1). _eprint: https://doi.org/10.1146/annurev-linguistics-031220-010811, 519–540. https://doi.org/10.1146/annurev-linguistics-031220-010811. https://doi.org/10.1146/annurev-linguistics-031220-010811 (29 June, 2023).

Denić, Milica, Shane Steinert-Threlkeld & Jakub Szymanik. 2022. Indefinite Pronouns Optimize the Simplicity/Informativeness Trade-Off. *Cognitive Science* 46(5). e13142. https://doi.org/10.1111/cogs.13142.

Denić, Milica & Jakub Szymanik. 2023. *Languages Optimize the Trade-off between Lexicon Size and Average Utterance Length: A Case Study of Numeral Systems*. https://lingbuzz.net/lingbuzz/006748.

Evans, Nicholas & Stephen C. Levinson. 2009. The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences* 32(5). Publisher: Universiteit of Amsterdam, 429–448. https://doi.org/10.1017/S0140525X0999094X.

Feldman, Jacob. 2000. Minimization of Boolean Complexity in Human Concept Learning. *Nature* 407. 630–633. https://doi.org/10.1038/35036586.

von Fintel, Kai & Sabine Iatridou. 2008. How to Say Ought in Foreign: The Composition of Weak Necessity Modals. In Jacqueline Guéron & Jacqueline Lecarme (eds.), *Time and Modality*, vol. 75 (Studies in Natural Language and Linguistic Theory), 115–141. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-1-4020-8354-9_6.

von Fintel, Kai & Lisa Matthewson. 2008. Universals in semantics. *The Linguistic Review* 25. 139–201. https://doi.org/10.1515/TLIR.2008.004.

Fodor, Jerry. 1975. *The language of thought*. New York: Crowell.

Frank, Michael C & Noah D Goodman. 2012. Predicting Pragmatic Reasoning in Language Games. *Science (New York, N.Y.)* 336. 998.

Gluckman, John & Margit Bowler. 2020. The expression of modality in Logoori. *Journal of African Languages and Linguistics* 41(2). 195–238. https://doi.org/doi:10.1515/jall-2020-2010. https://doi.org/10.1515/jall-2020-2010.

Goodman, Noah D, Joshua B Tenenbaum, Jacob Feldman & Thomas L Griffiths. 2008. A rational analysis of rule-based concept learning. *Cognitive Science* 32(1). 108–154. https://doi.org/10.1080/03640210701802071.

Gray, James. 2021. Variable Modality in Pintupi-Luritja Purposive Clauses. *Languages* 6(1). 52. https://doi.org/10.3390/languages6010052.

Guo, Qingxia, Nathaniel Imel & Shane Steinert-Threlkeld. 2022. A database for modal semantic typology. In *4th workshop on research in computational linguistic typology and multilingual NLP (SIGTYP2022)*, 42–51. Seattle, Washington: Association for Computational Linguistics. https://aclanthology.org/2022.sigtyp-1.6.

Hellwig, Birgit. 2011. *A grammar of goemai* (Mouton grammar library ; 51). Berlin ; Boston: De Gruyter Mouton.

Hockett, Charles F. 1960. The origin of speech. *Scientific American* 203. 88–96. http://www.isrl.uiuc.edu/~amag/langev/paper/charles60theOrigin.html.

Kemp, Charles & Terry Regier. 2012. Kinship categories across languages reflect general communicative principles. *Science (New York, N.Y.)* 336(6084). 1049–1054. https://doi.org/10.1126/science.1218811.

Kemp, Charles, Yang Xu & Terry Regier. 2018. Semantic typology and efficient communication. *Annual Review of Linguistics*. 1–23. https://doi.org/10.1146/annurev-linguistics-011817-045406.

Kratzer, Angelika. 1981. The Notional Category of Modality. In Hans-Jürgen Eikmeyer & Hannes Rieser (eds.), *Words, Worlds, and Context*, 38–74. Walter de Gruyter.

Levinson, Stephen C. 2012. Kinship and Human Thought. *Science (New York, N.Y.)* 336(6084). 988–989. https://doi.org/10.1126/science.1222691. https://www.science.org/doi/abs/10.1126/science.1222691.

Matthewson, Lisa. 2004. On the Methodology of Semantic Fieldwork. *International Journal of American Linguistics* 70(4). Publisher: The University of Chicago Press, 369–415. https://doi.org/10.1086/429207.

Matthewson, Lisa. 2013. Gitksan Modals. *International Journal of American Linguistics* 79(3). 349–394. https://doi.org/10.1086/670751.

Matthewson, Lisa. 2019. Modality. In Maria Aloni & Paul Dekker (eds.), *The Cambridge Handbook of Formal Semantics*, 525–559. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9781139236157.019.

Močnik, Maša & Rafael Abramovitz. 2019. A Variable-Force Variable-Flavor Attitude Verb in Koryak. In Julian J Schlöder, Dean McHugh & Floris Roelofsen (eds.), *22nd Amsterdam Colloquium*, 494–503.

Mollica, Francis, Geoff Bacon, Noga Zaslavsky, Yang Xu, Terry Regier & Charles Kemp. 2021. The forms and meanings of grammatical markers support efficient communication. *Proceedings of the National Academy of Sciences* 118(49). Publisher: Proceedings of the National Academy of Sciences, e2025993118. https://doi.org/10.1073/pnas.2025993118.

Nauze, Fabrice Dominique. 2008. *Modality in typological perspective*. Universiteit van Amsterdam phd.

Nikolaeva, Irina. 2014. *A grammar of tundra nenets* (Mouton grammar library ; Volume 65). Berlin, [Germany] ; Boston, [Massachusetts]: De Gruyter Mouton.

Piantadosi, Steven T, Joshua B Tenenbaum & Noah D Goodman. 2016. The logical primitives of thought: Empirical foundations for compositional cognitive models. *Psychological Review* 123(4). 392–424. https://doi.org/10.1037/a0039980.

van de Pol, Iris, Paul Lodder, Leendert van Maanen, Shane Steinert-Threlkeld & Jakub Szymanik. 2021. Quantifiers satisfying semantic universals are simpler. In *Cognitive Science Society (Cogsci-2021)*, vol. 43.

Pyatkin, Valentina, Shoval Sadde, Aynat Rubinstein, Paul Portner & Reut Tsarfaty. 2021. The possible, the plausible, and the desirable: Event-based modality detection for language processing. In *Annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing (volume 1: Long papers)*, 953–965. Online: Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.acl-long.77. https://aclanthology.org/2021.acl-long.77.

Quine, W. V. 1952. The problem of simplifying truth functions. *The American Mathematical Monthly* 59(8). Publisher: Mathematical Association of America, 521–531. http://www.jstor.org/stable/2308219.

Regier, Terry, Charles Kemp & Paul Kay. 2015. Word meanings across languages support efficient communication Informativeness and simplicity as competing principles. In Brian MacWhinney & William O'Grady (eds.), *The Handbook of Language Emergence*, 237–263. Wiley.

Rullmann, Hotze, Lisa Matthewson & Henry Davis. 2008. Modals as distributive indefinites. *Natural Language Semantics* 16(4). 317–357. https://doi.org/10.1007/s11050-008-9036-0.

Skyrms, Brian. 2010. *Signals: evolution, learning and information*. Oxford University Press.

Steinert-Threlkeld, Shane. 2020a. An Explanation of the Veridical Uniformity Universal. *Journal of Semantics* 37(1). 129–144. https://doi.org/10.1093/jos/ffz019.

Steinert-Threlkeld, Shane. 2020b. Quantifiers in natural language optimize the Simplicity/Informativeness trade-off. In Julian J Schlöder, Dean McHugh & Floris Roelofsen (eds.), *The 22nd Amsterdam Colloquium*, 513–522.

Steinert-Threlkeld, Shane. 2021. Quantifiers in Natural Language: Efficient Communication and Degrees of Semantic Universals. *Entropy. An International and Interdisciplinary Journal of Entropy and Information Studies* 23(10). 1335. https://doi.org/10.3390/e23101335.

Steinert-Threlkeld, Shane, Nathaniel Imel & Qingxia Guo. 2023. A Semantic Universal for Modality. *Semantics & Pragmatics* 16(1). https://doi.org/10.3765/sp.16.1.

Steinert-Threlkeld, Shane & Jakub Szymanik. 2019. Learnability and Semantic Universals. *Semantics & Pragmatics* 12(4). https://doi.org/10.3765/sp.12.4.

Steinert-Threlkeld, Shane & Jakub Szymanik. 2020. Ease of Learning Explains Semantic Universals. *Cognition* 195. https://doi.org/10.1016/j.cognition.2019.104076.

Uegaki, Wataru. 2021. The informativeness / complexity trade-off in the domain of Boolean connectives. *Linguistic Inquiry*.

Uegaki, Wataru & Ella Hannon. 2022. *Cross-linguistic dataset of force-flavour combinations in modal elements*.

Vander Klok, Jozina. 2013a. Pure Possibility and Pure Necessity Modals in Paciran Javanese. *Oceanic Linguistics* 52(2). 341–374. https://doi.org/10.1353/ol.2013.0017.

Vander Klok, Jozina. 2013b. Restrictions on semantic variation: A case study on modal system types. In *Workshop on Semantic Variation*.

Vander Klok, Jozina. 2021. Revised Modal Questionnaire for Cross-Linguistic Use.

Wiebe, Janyce, Theresa Wilson & Claire Cardie. 2005. Annotating Expressions of Opinions and Emotions in Language. *Language Resources and Evaluation* 39(2). 165–210. https://doi.org/10.1007/s10579-005-7880-9. https://doi.org/10.1007/s10579-005-7880-9.

Xu, Yang & Terry Regier. 2014. Numeral systems across languages support efficient communication: From approximate numerosity to recursion. In *Cognitive science society (cogsci-2014)*, 1802–1807.

Xu, Yang, Terry Regier & Barbara C. Malt. 2016. Historical Semantic Chaining and Efficient Communication: The Case of Container Names. *Cognitive Science* 40(8). 2081–2094. https://doi.org/10.1111/cogs.12312.

Yanovich, Igor. 2016. Old English *motan, variable-force modality, and the presupposition of inevitable actualization. *Language* 92(3). 489–521. https://doi.org/10.1353/lan.2016.0045.

Zaslavsky, Noga, Karee Garvin, Charles Kemp, Naftali Tishby & Terry Regier. 2022. The evolution of color naming reflects pressure for efficiency: Evidence from the recent past. *Journal of Language Evolution*. lzac001. https://doi.org/10.1093/jole/lzac001. https://doi.org/10.1093/jole/lzac001 (29 June, 2023).

Zaslavsky, Noga, Charles Kemp, Terry Regier & Naftali Tishby. 2018. Efficient compression in color naming and its evolution. *Proceedings of the National Academy of Sciences* 115(31). 7937–7942. https://doi.org/10.1073/pnas.1800521115.

Zipf, George Kingsley. 1949. *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology*. Addison-Wesley Press.

Appendices

A Estimating Minimum Descriptions in a Modal Language of Thought

A.1 A Boolean Algebra as the Modal Language of Thought

Our Language of Thought for modal meanings takes the form of a boolean algebra, where formulae are interpreted as sets of force-flavor pairs in the modal semantic space. Concretely, the semantic space is a set $\mathbb{M} = Q \times F$, where Q contains atoms for the possible modal forces (e.g. $\{w, s\}$ for *weak, strong*) and F contains atoms for the possible modal flavors (e.g. $\{e, d, c\}$ for *epistemic, deontic, circumstantial*). It is natural to think of the semantic space as a finite model for the standard language of propositional logic.

Since our representation language is interpreted with respect to this finite, two-dimensional model, the standard operators of disjunction $\vee$, conjunction $\wedge$ and negation $\neg$ behave slightly differently as they normally would, such that some laws hold in our algebra that do not correspond to sound inferences in standard propositional logic. These nonstandard laws work to our advantage as we search over possible minimizations of LoT formulae when estimating the minimum description length complexity of a modal meaning. The remainder of this appendix describes the laws (some that are specific to our specific algebra, and others that hold in boolean algebra generally) that we used as inferences for this complexity estimation.

Disjunction laws

(Commutativity) $x \vee y = y \vee x$

(Exhaustification 1) $\bigvee_{i=1}^{|Q|} q_i = \top$

(Exhaustification 2) $\bigvee_{i=1}^{|F|} f_i = \top$

(Identity) $\bot \vee x = x$

(Idempotence) $x \vee x = x$

Conjunction laws

(Commutativity) $x \wedge y = y$

(Annihilation 1) $\forall i \neq j : q_i \wedge q_j = \bot$

(Annihilation 2) $\forall i \neq j : f_i \wedge f_j = \bot$

(Idempotence) $x \wedge x = x$

Other laws

(Distributivity of $\vee$ over $\wedge$) $(x \wedge y) \vee (x \wedge z) = x \vee (y \wedge z)$

(Distributivity of $\wedge$ over $\vee$) $(x \vee y) \wedge (x \vee z) = x \wedge (y \vee z)$

(Complementation 1) $\forall q_i \in Q : \bigvee_{j \neq i} q_j = \neg q_i$

(Complementation 2) $\forall f_i \in F : \bigvee_{j \neq i} f_j = \neg f_i$

(DeMorgan) $\neg(x \wedge y) = \neg x \vee \neg y$

Example

Here we give an example of how the boolean algebra laws described above can be applied to map a modal meaning to its minimum-length formula. Each force-flavor pair (q_i, f_j) can be written as a conjunction $q_i \wedge f_j$, so that the entire set of meanings that the modal can express can be represented as a disjunction of these force-flavor conjunctions, resulting in a disjunctive normal form (DNF).

In our example, we consider a modal that can express every force-flavor pair combination except for the pair (q_1, f_3), and show how to derive the minimal formula $\neg(f_3 \wedge q_1)$ from this long DNF. In practice, one can obtain this formula immediately by converting the set of force-flavor points the modal *cannot* express to a DNF, and negating the resulting minimal formula, but for illustrative purposes we demonstrate the 'inefficient' derivation.

	f_1	f_2	f_3
q_1	✓	✓	
q_2	✓	✓	✓

$$\mapsto (q_1 \wedge f_1) \vee (q_1 \wedge f_2) \vee (q_2 \wedge f_1) \vee (q_2 \wedge f_2) \vee (q_2 \wedge f_3)$$

$$= q_1 \wedge (f_1 \vee f_2) \vee q_2 \wedge (f_1 \vee f_2 \vee f_3) \quad \text{(Distributivity)}$$

$$= q_1 \wedge (f_1 \vee f_2) \vee q_2 \wedge \top \quad \text{(Exhaustification 2)}$$

$$= q_1 \wedge (f_1 \vee f_2) \vee q_2 \quad \text{(Conjunctive Identity)}$$

$$= q_1 \wedge \neg f_3 \vee q_2 \quad \text{(Complementation 2)}$$

$$= \neg(\neg q_1 \vee f_3) \vee q_2 \quad \text{(DeMorgan)}$$

$$= \neg(q_2 \vee f_3) \vee q_2 \quad \text{(Complementation 1)}$$

$$= \neg((q_2 \vee f_3) \wedge \neg q_2) \quad \text{(DeMorgan)}$$

$$= \neg((q_2 \vee f_3) \wedge q_1) \quad \text{(Complementation 1)}$$

$$= \neg(q_1 \wedge q_2 \vee q_1 \wedge f_3) \quad \text{(Distributivity of } \wedge \text{ over } \vee\text{)}$$

$$= \neg(\bot \vee q_1 \wedge f_3) \quad \text{(Annihilation)}$$

$$= \neg(q_1 \wedge f_3) \quad \text{(Disjunctive Identity)}$$

A.2 Shortest Formula Heuristic Search Algorithm

Here we describe the heuristic algorithm (mentioned in Section 4.1) that we use to estimate the minimum description length of every modal meaning considered in our analysis. The algorithm works in two steps. First, it converts a modal meaning to a formula in the modal LoT described in Appendix A.1. Then, it performs a breadth-first search over possible minimizations of this formula, using the boolean algebra laws described in Appendix A.1 as logical inferences. These inferences are stored in `operations`, and are implemented as transformations of the syntactic trees of formulae in the LoT. Importantly, when exploring possible minimizations, our search algorithm does not require candidate formulae to be shorter than the current one. This allows our procedure to have access to solutions contained in search paths for which expression complexity is not generally decreasing. While we are confident that this procedure works well in practice in our experiments, we do not have provable guarantees that it finds the shortest formulae for meanings in larger semantic spaces; hence, this algorithm is merely a heuristic.

Algorithm 1 Approximating the shortest formula for a modal.

```
function SHORTEN(m, operations, max_iterations)   ▷ a modal m is a set of meaning points
    f ← points_to_dnf(m)                          ▷ f is a DNF formula
    to_visit ← {f}
    i ← 0
    shortest ← f
    while to_visit ≠ ∅ do                         ▷ breadth-first search
        if i = max_iterations then
            break
        end if
        next ← to_visit.dequeue()
        children = ∅
        for operation in operations do
            children.enqueue(operation(next))
        end for
        for child in children do
            if child ≠ next then
                to_visit.enqueue(child)
            end if
        end for
        i ← i + 1
        if complexity(next) < complexity(shortest) then
            shortest ← next
        end if
    end while
    return shortest
end function
```

Nathaniel Imel
Social Science Plaza B, 2nd Floor
Irvine, California 92697
nimel@uci.edu

Qingxia Guo
Guggenheim Hall, 4th Floor
Seattle, Washington 98195
qg07@uw.edu

Shane Steinert-Threlkeld
Guggenheim Hall, 4th Floor
Seattle, Washington 98195
shanest@uw.edu