New Logical Perspectives on Monotonicity

Johan van Benthem[1] and Fenrong Liu[2]

[1] University of Amsterdam, Stanford University, and Tsinghua University
johan@stanford.edu
[2] Tsinghua University and University of Amsterdam
fenrong@tsinghua.edu.cn

Abstract. Monotonicity-based inference is a fundamental notion in the logical semantics of natural language, and also in logic in general. Starting in generalized quantifier theory, we distinguish three senses of the notion, study their relations, and use these to connect monotonicity to logics of model change. At the end we return to natural language and consider monotonicity inference in linguistic settings with vocabulary for various forms of change. While we mostly raise issues in this paper, we do make a number of new observations backing up our distinctions.

Keywords: monotonicity · generalized quantifiers · model change.

1 Varieties of monotonicity for generalized quantifiers

Basic patterns. Monotonicity is a property that is used extensively in linguistics and logic. Many valid reasoning patterns involve monotonicity, in particular with sentences containing generalized quantifiers. Here are four possible cases with a binary generalized quantifier Q and two predicate arguments A and B:

↑MON $Q(A, B)$ and $A \subseteq C$, then $Q(C, B)$
↓MON $Q(A, B)$ and $C \subseteq A$, then $Q(C, B)$
MON↑ $Q(A, B)$ and $B \subseteq C$, then $Q(A, C)$
MON↓ $Q(A, B)$ and $C \subseteq B$, then $Q(A, C)$

For instance, the universal quantifier "all" is downward monotonic in its left argument and upward in its right argument, thus exemplifying the type ↓MON↑. If we want to stress possible dependence of the quantifier on a total domain of discourse D, the binary notation $Q(A, B)$ will be extended to a ternary $Q_D(A, B)$.[3]

Three senses. While the preceding definitions seem clear, intuitive explanations of monotonicity inference in natural language sometimes appeal to slightly, but subtly different notions. This note identifies three possible interpretations, and then goes on to discuss these in a variety of logical settings, raising new issues in the process. We will focus on upward monotonicity in what follows, though our analysis also applies to downward monotonicity.

[3] An extensive overview of monotonicity inference with generalized quantifiers can be found in (Peters and Westerståhl, 2006).

To introduce what we have in mind, consider the following three examples:

(1a) Some boys dance. (1b) Some people dance.

The upward monotonic step from (1a) to (1b) may be called *Predicate Replacement* in the same domain of objects. The more specific (stronger) predicate "boys" (A) is replaced by the more general (weaker) predicate "people" (C).

Next, consider a case that feels intuitively different, where the *same* predicate changes its extension. For a long time, whales were thought of as fish, but then it was found they are mammals, and the range of "mammal" was extended.

(2a) Some mammals (excluding whales) live over a hundred years.

(2b) Some mammals (including whales) live over a hundred years.

Here, a predicate acquires more members in the same domain. Call this view of monotonicity *Extension Increase*. At this stage, however, we have a distinction without a difference. Predicate Replacement and Extension Increase are the same for quantifiers viewed as set relations. But as we will see later, the distinction starts making sense when we have both syntax and semantics. [4]

But there is yet a stronger form of monotonicity, where the domain itself can be enlarged. Suppose that we are first talking about Asians, and next about all people in the World. The following monotonicity inference is valid:

(3a) Some musicians are Chinese (in Asia).

(3b) Some musicians are Chinese (in the whole World).

Let us call this form of monotonicity *Domain Enlargement*. The predicate "musician" does not change its extension in the old Asian domain, but we now consider its full extension in the new World domain. Of course, since "some" satisfies both Extension Increase and Domain Enlargement, we can even combine the two. The resulting Enlargement Monotonicity is illustrated in the following diagram:

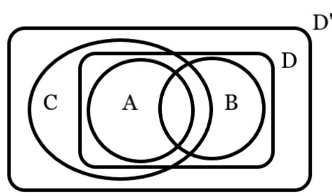

While the distinction between keeping the domain fixed or extending it for monotonicity seems intuitive, it, too, collapses – when we accept an assumption called *Extension* that is commonly made for generalized quantifiers:

EXT if $A, B \subseteq D \subseteq D'$, then $Q_D(A, B)$ iff $Q_{D'}(A, B)$

Fact 1 *With* EXT*, upward monotonic Predicate Replacement (I) and Enlargement Monotonicity (II) are equivalent conditions on quantifiers.*

[4] There is also an intuitive temporal aspect to the whales example, where extensions change with the passage of time. Such more intensional aspects of monotonicity inference will be considered briefly at the end of this paper.

Proof. We only consider upward monotonicity in the left argument. *From (II) to (I).* Predicate replacement is clearly a special case of Enlargement Monotonicity when the domain does not change. *From (I) to (II).* Let $Q_D(A, B)$ and $A \subseteq C$, and $C, D \subseteq D'$. By EXT, we have $Q_{D'}(A, B)$, and then by (I), $Q_{D'}(C, B)$. □

Shrinking domains with downward monotonic inference gives a similar result. However, EXT is crucial to all of this.

Doing without EXT. With quantifiers whose meaning involves the domain D in an essential way, monotonicity becomes a much richer notion, where Conservativity is no longer a prominent constraint. We will only illustrate this here, since these quantifiers seem much less studied. Consider the quantifier *"many A are B"* in one plausible sense of relative proportion:

$$\frac{| A \cap B |}{| A |} > \frac{| B |}{| D |}$$

where D is the whole domain.[5] It is illustrated in the following diagram, where the numbers of objects in the different zones have been marked by x, m, k, n:

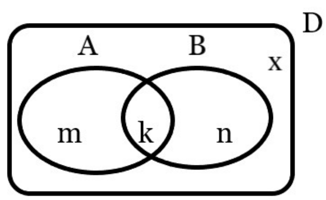

According to the above definition, we have

$$\tfrac{k}{m+k} > \tfrac{k+n}{m+n+k+x}, \quad \text{or equivalently,} \quad kx > mn$$

Now our earlier distinction between keeping the domain fixed, or extending it makes sense. Also, the notion of monotonicity acquires new options.

Clearly, MON↑ in its standard sense can fail when we enlarge B with objects in D outside of A that increase the frequency of B in D, but not in A. However, the new setting allows for more subtle forms of monotonicity.

Here is a natural candidate, keeping the A's fixed. If we enlarge B *inside of A only*, we regain MON↑. To illustrate why, just add one object of A to B, raising k to $k + 1$ while lowering m to $m - 1$. Then we have

$$\tfrac{k+1}{m-1+k+1} = (\tfrac{k}{m+k} + \tfrac{1}{m+k}) > \tfrac{k+1+n}{m+n+k+x}$$

Next, consider domain extension. Clearly just increasing B can make *"many A are B"* false. But we can also enlarge just $A \cap B$, putting a new B-object s in the new A. This time, *"many A are B"* remains true since it implies

[5] In particular, with this definition, it is never true that D-relatively many A are D. We will not discuss other variants of relative "many" here.

$\frac{k+1}{m+k+1} > \frac{k+1+n}{m+k+1+n+x}$[6]

Merging logic and counting. There is no general theory of types of monotonicity in this extended setting for quantificational reasoning. Note that monotonicity as discussed here fits naturally with qualitative perspectives on numerical formulas with addition, multiplication and other elementary operations, giving us global information about how functions grow as argument values change.[7] Thus, the right format for this broader setting may be a system of 'counting logic' mixing set-theoretic and arithmetical components. This would fit with the intuitive idea that quantifiers are at heart about counting, so that actual reasoning with quantifiers may well be a mix of just this kind.[8]

Monotonicity calculus. In practice, upward and downward monotonic inferences are equally important. Syntactically, these are triggered by *positive* and *negative* ocurrences, respectively, of the predicate replaced in the inference. And since quantifiers can occur embedded in further linguistic constructions, a calculus is needed for computing positive and negative occurrences of predicates inside complex expressions. For instance, in "every pot has a lid", "pot" is negative, supporting a downward inference, while the embedded "lid" is positive, supporting an upward inference. Taken together, it follows, e.g., that "every iron pot has a cover". A precise Monotonicity Calculus keeping track of positive and negative syntactic occurrences can be stated in terms of a categorial grammar for constructing complex expressions, cf. (van Benthem, 1991). While details of this system are not relevant to us here, its existence suggests looking at logical systems that contain quantifiers to take our analysis a step further.

2 Monotonicity in first-order logic

Two senses revisited. In first-order logic, a pilot system for a mathematical theory of generalized quantifiers, truth values of formulas depend on domains of models. In other words, EXT no longer holds when first-order syntax for quantifiers is taken into account. Two notions of monotonicity may be distinguished, where again we focus on the upward case to simplify the exposition:

(Mon-inf) From $\varphi(P)$ and $\forall x(P(x) \to Q(x))$, it follows that $\varphi(Q/P)$, where $\varphi(Q/P)$ is the result of replacing each occurrence of P in φ by Q.

[6] This inequality is equivalent to $kx+x > mn$ which is implied by the earlier $kx > mn$.

[7] A realistic concrete use of monotonicity in mathematics is the convergence test for improper integrals discussed in (Icard, Moss and Tune, 2017).

[8] In this combined calculus, monotonicity applies to both *set inclusion* for denotations and *greater-than* for numbers. The former is a type-lifting of the latter, and many more complex type-theoretic lifts support monotonicity reasoning (van Benthem, 1991). However, beyond these, in natural language monotonicity can apply to many orderings that are sui generis: conceptual, temporal, spatial, and so on. Can the style of analysis in this paper be generalized to cover these?

(Mon-sem) If $M, s \models \varphi(P)$ and $M \equiv^+_P M'$ (i.e., M and M' are the same model except for the interpretation of P, and $I(P) \subseteq I'(P)$), then $M', s \models \varphi(P)$.

These correspond to the earlier Predicate Replacement and Extension Increase.

Fact 2 *Inferential monotonicity is equivalent to semantic monotonicity.*

Proof. From (Mon-sem) to (Mon-inf). Suppose, for any model M and assignment s, that $M, s \models \varphi(P)$ and $M, s \models \forall x(P(x) \to Q(x))$. Now define a new model M' which is like M except that $I'(P) = [[Q]]^M$. Clearly $M \equiv^+_P M'$, so by Mon-sem, we have $M', s \models \varphi(P)$. By one direction of the standard Predicate Substitution Lemma for first-order logic, it then follows that $M, s \models \varphi(Q/P)$.

From (Mon-inf) to (Mon-sem). Suppose that $M, s \models \varphi(P)$ and $M \equiv^+_P M'$. Take a new predicate letter Q not occurring in $\varphi(P)$, and set $I(Q) = I'(Q) = [[P]]^{M'}$. Then in the model M, s, the two conditions for Mon-inf are satisfied, and therefore, $\varphi(Q/P)$ is true in M, s. But this implies, by the converse direction of the Predicate Substitution Lemma, that $M', s \models \varphi(P)$. □

The second half of this proof requires the availability of fresh predicates. We suspect that the above equivalence fails for first-order logic with a finite vocabulary while it still holds for subsystems such as monadic FOL.

For the earlier third sense of Domain Enlargement, see Section 3 below.

Single vs. multiple occurrences. In actual inferences based on Mon-inf, it is natural to focus on a single occurrence of the predicate P. Typically, this upward form is licensed when this occurrence of P is syntactically positive in φ. But note that the same P may also have negative occurrence in φ. For instance, in $P \wedge \neg(P \wedge Q)$, the first occurence of P is positive, supporting a MON↑ inference, but the second occurrence is negative, supporting a MON↓ inference.[9] However, our discussion also covers inferences with multiple replacements.

Interpolation and monotonicity calculus. Semantic monotonicity jumps from one model to another along the relation $\equiv^+_P$. A related general notion of transfer between models is this: φ *entails* ψ *along* R if, whenever MRN and $M \models \varphi$, then $N \models \psi$. This notion was introduced in (Barwise and van Benthem, 1999), which also proves the following version of Lyndon's Theorem for FOL:

Fact 3 *The following statements are equivalent for first-order formulas φ, ψ: (a) φ entails ψ along $\equiv^+_P$, (b) there exists a formula α containing only positive syntactic occurrences of P such that $\varphi \models \alpha \models \psi$.*

The required formulas α are generated by the grammar

$$P\mathbf{x} \mid (\neg)Q\mathbf{x}\,(Q \neq P) \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \exists x\varphi \mid \forall x\varphi.$$

Fact 3 can be seen as a completeness result for the monotonicity calculus of first-order logic. But to make this apply to generalized quantifier theory, one

[9] Many inferences are intuitively about single occurrences of parts of expressions. But some require comparing coordinated occurrences, like in the logical rule of Contraction, where two identical premises can be contracted to just one.

needs similar results for the logics $FOL(\mathbf{Q})$ consisting of first-order logic with added generalized quantifiers. This has been done in (Makovsky and Tulipani, 1977), using suitable extensions of the basic model-theoretic notions for FOL.[10]

Semantics that fit monotonicity inference. Here is another way of phrasing the preceding completeness issue. The monotonicity calculus is a proof system for practical reasoning. Is there a natural semantics for which it is complete? Interesting answers have been given, cf. the proposals considered in (Icard, Moss and Tune, 2017). In addition, here is a straightforward modal perspective.

In *modal state semantics* for first-order logic, the variable assignments of Tarski semantics are viewed as abstract states, and quantifiers $\exists x$ are then interpreted using arbitrary accessibility relations R_x between states. The result of this widening of standard models is a decidable modal sublogic of FOL which blocks all valid first-order consequences except for monotonicity and aggregation of universal statements under conjunction. To block the latter, a well-known move in modal logic is a step from binary accessibility relations between states to *neighborhood models* with state-to-set neighborhood relations. A straightforward neighborhood generalization of state models for FOL will validate essentially just the monotonicity inferences. For further details, and connections to generalized quantifiers, cf. (Andréka, van Benthem and Németi, 2017).

3 Logics for monotonicity-related model change

Intuitively, the third and second sense of monotonicity in Section 1 involve model change. In recent years, families of logics have been studied that analyze the effects of changing models, for the purposes of information update, world change, game play, or other concrete scenarios. These logics can code our earlier reasoning about monotonicity, while at the same time, they extend practical monotonicity inference to new settings. In this section we discuss some connections.

Predicate extension modalities. For an illustration, take the case of upward predicate monotonicity, and add the following modality to first-order logic[11]

$\langle \equiv^+_P \rangle \varphi$ for: φ is true in some $\equiv^+_P$-extension of the current model.

With the dual universal modality in the language, upward semantic monotonicity can now be formulated as an object-level validity of the system:

$$\varphi(P) \to [\equiv^+_P]\varphi(P)$$

As another example, the fact that positive occurrence of P in φ implies upward monotonicity is expressible by a set of valid implications in this language.

The new modality is very powerful, as it can express existential second-order quantifiers. To see this, take a first-order sentence defining discrete linear orders with a beginning but no end, (i). Next, with a unary predicate P, the formula

[10] Extensions to richer type logics of relevance to natural language seem an open problem, cf. (van Benthem, 1991) on the case of the Boolean Lambda Calculus.

[11] This device has not been studied yet in the literature, to the best of our knowledge, but as we shall see momentarily, it is close to second-order logic.

$$\forall x \neg Px \wedge \neg\langle \equiv^+_P \rangle(\exists x Px \wedge \forall x(Px \to \exists y(y < x \wedge Py)))$$

says there is no non-empty subset of the domain without a minimal element, (ii). The conjunction of (i) and (ii) defines the standard natural numbers, whose complete predicate logic (in a rich enough vocabulary) is non-arithmetical.

The expressive power of the monotonicity modality can be much less on fragments of FOL, representing more elementary settings for monotonicity reasoning.

Fact 4 *Adding* $\langle \equiv^+_P \rangle$ *to* monadic FOL *adds no expressive power at all.*

Proof sketch. The proof is by a syntactic normal form argument in the style of (van Benthem, Mierzewski and Zaffora Blando, 2020). Each monadic first-order formula is equivalent to a disjunction of the following form:

(i) global state descriptions that list which of the 2^k possible true/false combinations for k unary predicates are exemplified in the model, conjoined with (ii) local state descriptions for a finite set of variables.

Prefixing a modality $\langle \equiv^+_P \rangle$ distributes over the initial disjunction, and we are left with the modality over the described conjunctions. With this complete explicit syntactic description available, it is easy to read off what is expressed in terms of conditions that can be formulated entirely in monadic FOL.

Instead of an algorithm for deriving these conditions, we give an example:

$\langle \equiv^+_P \rangle(\neg\exists x(Px \wedge Qx) \wedge \exists x(Px \wedge \neg Qx) \wedge \exists x(\neg Px \wedge Qx) \wedge \neg\exists x(\neg Px \wedge \neg Qx) \wedge Px \wedge \neg Qx)$ is equivalent with the monadic formula $\neg\exists x(Px \wedge Qx) \wedge \exists x(\neg Px \wedge Qx) \wedge \exists x((Px \wedge \neg Qx) \vee (\neg Px \wedge \neg Qx)) \wedge Px \wedge \neg Qx$[12]

A similar closure argument will work for monadic first-order logic with identity.

However, adding the monotonicity modality to another weak decidable fragment of FOL already yields much higher complexity. The modal 'fact change logic' of (Thompson, 2019) adds a modality $\langle +p \rangle\varphi$ to basic modal logic saying that making p true in the current world makes φ true there. Under the standard translation of modal logic into first-order logic, this becomes a fragment of the language of FOL plus a special case of the modality $\langle \equiv^+_P \rangle$. Fact change logic is still axiomatizable, but unlike the basic modal logic, it is undecidable.

Domain enlargement. The third sense of monotonicity involved Domain Enlargement. This suggests adding a modality $\langle \subseteq \rangle\varphi$ to FOL saying that φ is true in some extension of the current model.[13] This logic encodes the usual facts such as preservation of existential first-order formulas under model extensions. But again, this system in general has very high complexity. For instance, it can define that a first-order formula φ is satisfiable, by taking a fresh unary predicate letter P not interpreted in the current model, and stating that φ can be made true relativized to P: $\langle \subseteq \rangle(\varphi)^P$. As before, fragments are better behaved, and of particular interest are stepwise addition (or deletion) of objects in a current model,

[12] This can be simplified to $\neg\exists x(Px \wedge Qx) \wedge \exists x(\neg Px \wedge Qx) \wedge \exists x \neg Qx \wedge Px \wedge \neg Qx$.
[13] Enlargement Monotonicity is then expressed by modal combinations like $\langle \subseteq \rangle\langle \equiv^+_P \rangle$.

(Renardel de Lavalette, 2001), in line with intuitive reasoning about diagrams with generalized quantifiers. We do not pursue this topic here.

Information update meets monotonicity inference. A final setting for model change lets inference steps meet with semantic information updates, a natural combination in practical problem solving (van Benthem, 2011). For a concrete setting, in 'public announcement logic' (PAL), modalities $[!\varphi]\psi$ express that ψ will be true at the current world after original model has been updated with the information that φ is true. For details on the logic PAL, see (Baltag and Renne, 2016). What upward monotonicity inferences are allowed here?

There are two places where these inferences can occur. First it is easy to see that the 'postcondition' ψ of formulas $[!\varphi]\psi$ allows for standard monotonic inference to $[!\varphi](\psi \vee \alpha)$, and similar weakenings are allowed for positive occurrences of p in ψ that are not in the scope of dynamic modalities contained in ψ.

But with p inside the announced φ, things are more complicated. $[!\varphi]\psi$ does not imply $[!(\varphi \vee \alpha)]\psi$: such a monotonic replacement may give *weaker information*, true in more worlds, changing the original update to a larger submodel where earlier effects can be blocked. For instance, for atomic facts p, the formula $[!p]Kp$ is valid in PAL: after receiving the information that p an agent will know that p. However, the formula $[!(p \vee q)]Kp$ with a weaker announcement is obviously not valid. In contrast, monotonicity in the postcondition does tell us that from stronger announced content weaker facts can become known. For instance, $[!(p \wedge q)]Kq$ is valid: we can also learn parts of what was announced.[14]

Dynamic monotonicity. But actually, a more dynamic form of monotonicity inference may be natural in the PAL environment, triggered by a dynamic take on inclusion viewed as a relation between informational actions. Let us say that an announcement (not a proposition) $!\varphi$ *entails* an announcement $!\psi$ if

the implication $[!\varphi]\alpha \leftrightarrow [!\varphi][!\psi]\alpha$ is valid in PAL for all formulas α.

One can think of this in Gricean terms, where stating $!\psi$ after $!\varphi$ would not be appropriate, as it adds no information. Viewed as an inclusion of actions, this sort of connection can trigger inferences. The logic PAL contains information about what can be deduced from entailments between announcements.[15] This is just one way of thinking. There are other natural notions of dynamic entailment – but we must leave the study of dynamic monotonicity to another occasion.

All this leads to a question. A *Lyndon-style preservation theorem* capturing semantic monotonicity in PAL formulas in syntactic terms remains to be found.

[14] It is easy to see with simple concrete examples of PAL update that downward monotonicity fails as well for announced formulas: $[!\varphi]\psi$ does not imply $[!(\varphi \wedge \alpha)]\psi$.

[15] The exact information content of an announcement $!\varphi$ is that φ *was* true before the announcement (the caveat is needed since announcing an epistemic statement φ might change its truth value), and if ψ subsequently adds no new information, this means that the $!\psi$ update does not change the model. Thus, a way of taking dynamic entailment is as a valid implication $Y\varphi \to \psi$, where Y is a one-step backward-looking temporal operator beyond the language of PAL, cf. (Sack, 2007).

However, this is not yet a precise question. To understand what might be involved here, note that moving to a larger submodel through a weaker update does preserve some earlier postconditions ψ, namely those that are *existentially definable*. Thus, a Lyndon result in the dynamic PAL setting may have to simultaneously analyze monotonicity and preservation under model extensions. Also, since we are in an intensional setting with formulas referring to different models, the inclusion triggers for monotonicity inferences need some care. Just inclusion in an initial model need not suffice for justifying replacement in postconditions referring to updated models: we must have triggers of the right strength, or in semantic terms: inclusion of denotations in all relevant models. [16]

4 Back to natural language

Dynamic logics for model change are useful tools for formalizing the metatheory of monotonicity and much else besides. But they can also model concrete inferences in a setting of instructions for change. In this final section, we briefly list some possible repercussions of the preceding technical topics when we return to generalized quantifiers in natural language, the area we started with.

Linguistics expressions of change. Descriptions of changes in the world or instructions for achieving these changes occur explicitly in natural language. For instance, the dynamic modality of public announcement logic suggests analogies with the verb *"to learn"*, which describes a change in information state. The earlier technical observations about PAL then suggest linguistic questions about inferences that go with learning. If we view "learn that A" as a description of what the agent comes to know after the learning, A is a postcondition that allows the upward monotonic conclusion "learn that $A \vee B$". But if we take the A to be the content of the message leading to the learning, we are rather talking about an announcement $!A$ where upward inference is not allowed, or at least tricky.

Many action verbs deserve attention here, such as "change", "make", or, closer to our second and third senses of monotonicity: "add", "increase", or "remove". As an example, consider whether the following inference is valid:

(4a) All A are B.

(4b) Increasing the A's is increasing the B's.

Here we see an ambiguity that matches our discussion of various senses of upward monotonicity in Section 1. If we increase only the extension of A in some fixed domain, then B might stay the same. But if we add a new object that is A and

[16] To make the above questions fully precise, we need to define syntactic polarity of occurrences in PAL formulas, where occurrences inside announced formulas may lack polarity. Also, given the intensional setting for PAL of a universe of many epistemic models connected through updates, the earlier semantic notion of monotonicity can be phrased in a number of ways. Finally, we need not confine ourselves to syntactic properties of single occurrences of predicates. A proper notion of monotonic inference for formulas $[!\varphi]\alpha$ might involve correlated *simultaneous* replacements of proposition letters in both φ and α. We leave these detailed issues for follow-up work.

insist on the premise, then indeed, we have also increased the number of B's.[17] So, there are options for taking proposed inferences in a dynamic setting.[18]

Also, the status of the inclusion premise needs attention. We demonstrate this with our next example. Perhaps most centrally, while classical monotonicity inference focuses on what *is* the case, the dynamic counterpart verb is *"become"*. Inferences with all of these expressions seem to involve intensional phenomena.

Monotonicity inference and intensionality. Consider this inference:

(5a) Prime ministers of India are male.

Indira Gandhi became PM of India.

(5b) Indira Gandhi became male.

This is obviously incorrect. Indira Gandhi's election *falsified* the generalization expressed in the first premise. The point is that the premise is sensitive to moments in time, and can change its truth value as events happen.

We are in familiar more general territory now, monotonicity inference in intensional contexts and modal logics. These generally require modified inclusion statements, modalized to the right degree. Something that would work in all cases is a modalized "strong inclusion" true in all worlds, but the inclusion may also be more specific to the intended conclusion. If prime ministers of India were granted legal emergency powers just before Indira Gandhi's election, then we would be justified in concluding that she acquired such powers, even if that inclusion was not always the case in history. For more on monotonicity inference in the setting of modal logic, we refer to (Aloni, 2005) and (Yan and Liu, 2020).[19][20]

These two brief examples may have shown how technical dynamic logics of change connect naturally with linguistic phenomena, in particular, the monotonicity inferences long studied in formal semantics. Once we take this perspective, many further connections suggest themselves. Here is a last illustration.

[17] (Sun and Liu, 2020) discuss such inference patterns in the ancient Chinese language.

[18] With this richer linguistic vocabulary in monotonicity reasoning, the more general orderings of Footnote 7 may also come to the fore. Thomas Icard (p.c.) gives the nice example of "The tree is tall. The tree grows. Therefore, the tree is still tall."

[19] The difference between inclusions locally true in the actual world and inclusions true also in other worlds remains somewhat hidden in common phrasings of upward monotonicity inference as a pattern "from $\varphi(P)$ to $\varphi(P \vee Q)$". The inclusion from P to $P \vee Q$ is universally valid, so usable anywhere.

[20] There are many further intensional aspect to monotonicity inference that we cannot address here. For instance, such inferences seem sensitive to *description*. In the ancient Mohist example that "Your sister is a woman. But loving your sister is not loving a woman", the issue may be under which description we are viewing the loving ('as a relative' vs. 'romantically'). This distinction is widespread. Oedipus killed a man on the road, but did not realize that the man was his father. Did he kill his father? Under one description: yes, under another: no. For many further instances of the role of description in intensional contexts, see (Aloni, 2001), (Holliday and Perry, 2015). Should we consider a more refined notion of monotonicity inference where inference can take place at either the level of denotations, or that of descriptions?

Monotonicity inference as topic dynamics. In line with dynamic views of natural language use, we can also view drawing an inference itself as a dynamic activity (van Benthem, 2011). A conclusion is often not something that just passively 'follows' (from) the premises. In addition, it can also be an active means of *changing*, or at least modifying the *topic* of discussion or investigation. In this sense, a monotonicity inference from p to $p \vee q$ is not just a 'weakening', or a form of non-relevant reasoning to be banned, but the introduction of a new topic. Indeed, topic change is again a general phenomenon for which dynamic modal logics exist, so then we have closed a circle in our considerations.

5 Conclusion

We have identified three different intuitive senses of monotonicity inference. In standard generalized quantifier theory these largely amount to the same thing. However, once we drop the usual GQT assumption of Extension, differences between the various senses emerge, including new forms of monotonicity. These came out clearly in systems that describe counting and logical inference on a par. After all, intuitively, quantifiers seem a place where logic meets quantitative reasoning. Next, when embedding quantifiers in richer languages, our three senses came apart in classical first-order logic, and yielded a number of interesting issues, including interpolation and completeness for generalized semantics. Going to less familiar settings, monotonicity also connected in interesting ways with new (modal) logics of model change, leading to an array of new questions. Finally, we have suggested that all this technical development may be taken back to natural language, suggesting a fresh look at the interplay of monotonicity inference with the rich linguistic vocabulary for expressing change.

Acknowledgements We thank Thomas Icard, Zhiqiang Sun, Jouko Väänänen, Dag Westerståhl, Jialiang Yan, the audience of the DALI Workshop (Prague 2020) and the referees of successive versions of this paper for helpful comments and useful information. This research is supported by the Major Program of the National Social Science Foundation of China (NO. 17ZDA026).

References

1. M. Aloni, *Quantification under Conceptual Covers*. PhD thesis, University of Amsterdam, Amsterdam, 2001
2. M. Aloni, Individual Concepts in Modal Predicate Logic, *Journal of Philosophical Logic*, vol. 34, pp.1–64, 2005
3. H. Andréka, J. van Benthem and I. Németi, On a New Semantics for First-Order Predicate Logic, *Journal of Philosophical Logic*, 46(3), pp. 259–267, 2017
4. A. Baltag and B. Renne, Dynamic Epistemic Logic, Edward N. Zalta ed. *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition)
5. J. Barwise and J. van Benthem, Interpolation, Preservation, and Pebble Games.*Journal of Symbolic Logic*, 64 (2):881-903, 1999

6. J. van Benthem, K. Mierzewski and F. Zaffora Blando, The Modal Logic of Stepwise Removal, *Review of Symbolic Logic*, to appear
7. J. van Benthem, *Language in Action: Categories, Lambdas and Dynamic Logic.* North-Holland, Amsterdam, Studies in Logic and the Foundations of Mathematics, vol. 130, 1991
8. J. van Benthem, *Logical Dynamics of Information and Interaction.* Cambridge University Press, Cambridge, 2011
9. W. Holliday and J. Perry, Roles, Rigidity and Quantification in Epistemic Logic. In *Johan van Benthem on Logic and Information Dynamics*, Springer, Dordrecht, pp. 591-629, 2015
10. T. Icard, L. Moss and W. Tune, A Monotonicity Calculus and Its Completeness, in *Proceedings of the 15th Meeting on the Mathematics of Language*, 2017
11. J. A. Makowsky and S. Tulipani, Some Model Theory for Monotone Quantifiers.*Arch. math. Logik* 18, pp.115-134, 1977
12. S. Peters and D. Westerståhl, *Quantifiers in Language and Logic.* Clarendon Press, Oxford, 2006
13. G. Renardel de Lavalette, A Logic of Modification and Creation. In *Logical Perspectives on Language and Information*, Stanford, CSLI Publications, 2001.
14. J. Sack, Adding Temporal Logic to Dynamic-Epistemic Logic, Dissertation, University of Indiana, Bloomington, 2007.
15. Z. Sun and F. Liu, The Inference Pattern *Mou* Seen from Monotonic Reasoning, manuscript, Department of Philosophy, Tsinghua University
16. D. Thompson, Local Fact Change Logic, in F. Liu, H. Ono, Hiroakira, J.Yu eds. *Knowledge, Proof and Dynamics: The Fourth Asian Workshop on Philosophical Logic*, pp.73–96. Logic in Asia: Studia Logica Library, Springer Singapore, 2020
17. J. Yan and F. Liu, Monotonic Opaqueness in Deontic Contexts, in *Proceedings of the 5th Asian Workshop on Philosophical Logic*, Springer Singapore, 2020