

Meaning and Demonstration

Matthew Stone · Una Stojnic

Received: 16 October 2013 / Accepted: 17 September 2014 / Published online: 14 November 2014
© Springer Science+Business Media Dordrecht 2014

Abstract In demonstration, speakers use real-world activity both for its practical effects and to help make their points. The demonstrations of origami mathematics, for example, reconfigure pieces of paper by folding, while simultaneously allowing their author to signal geometric inferences. Demonstration challenges us to explain how practical actions can get such precise significance and how this meaning compares with that of other representations. In this paper, we propose an explanation inspired by David Lewis's characterizations of coordination and scorekeeping in conversation. In particular, we argue that words, gestures, diagrams and demonstrations can function together as integrated ensembles that contribute to conversation, because interlocutors use them in parallel ways to coordinate updates to the conversational record.

1 Introduction

In a famous video on YouTube,¹ Vi Hart presents a masterful origami proof of the Pythagorean theorem. With a judicious and precise series of folds (and a couple of rips), Hart creates a visualization of the areas involved in the Pythagorean theorem

¹<http://www.youtube.com/watch?v=z6lL83w131E>

M. Stone (✉)
Computer Science & Cognitive Science,
Rutgers University, Piscataway, NJ 08854 USA
e-mail: Matthew.Stone@Rutgers.Edu

U. Stojnic
Philosophy & Cognitive Science,
Rutgers University, New Brunswick, NJ 08901 USA
e-mail: Una.Stojnic@Rutgers.Edu

in a way that's both mathematically precise and makes the equivalence of these areas visually obvious. Hart's folds are examples of *demonstrations*—practical actions used by their authors to communicate precise ideas.² We'll be drawing quite closely on Hart's utterances and demonstrations as examples throughout the rest of the paper. We encourage you to look at the video now, so Hart's work is fresh in your mind; it is about three minutes long.

Demonstrations are extremely common, both in abstract reasoning and in collaborative activity. However, as Hart herself observes at the start of her video, demonstrations constitute a direct challenge to our typical preconceptions about reasoning and communication.

- 1 You don't need numbers or fancy equations to prove the Pythagorean theorem. All you need is a piece of paper. There's a ton of ways to prove it, and people are inventing new ones all the time, but I'm going to show you my favorite—except instead of looking at diagrams we're gonna fold it.

As we shall see in Section 2 of this paper, Hart's claim in (1) is exactly correct: Her demonstrations constitute a proof. Folds in origami are amenable to precise mathematical analysis, on a par with the mathematical analyses of drawing with straight edge and compass in Euclidean geometry (Alperin 2000; Lang 2010). The formal operations of origami mathematics provide a construction that corresponds precisely to the argument Hart offers; this construction establishes the truth of the Pythagorean theorem. Hart's folds concretize this formal construction in much the same way inscriptions on paper can concretize the constructions of Euclidean geometry. This is how she realizes her plan, as she describes it in (1), to "fold a proof". Indeed, we will suggest that the way Hart narrates her argument is sufficiently detailed and precise to establish mathematical truths, even for viewers who haven't previously entertained the abstractions of origami mathematics. Hart seems to target such viewers especially.

Though we concur with Hart that origami is mathematics, we will also argue that Hart is right to emphasize the difference between folding and the drawing of more conventional diagrams. Hart's demonstrations don't just create inscriptions for us to look at. They involve dynamic physical manipulations that remake the configuration of the paper in diverse ways over time. This is indispensable to their communicative effects, and an important part of their appeal. Since Hart's argument is inscribed directly into the changing world, recovering her argument seems to require viewers to recognize and understand what they see happening.

Hart's demonstrations thus manage to scaffold an abstract mathematical construction over our everyday understanding of the world. More precisely, with any of Hart's demonstrations, there is a kind of overlap between its practical effects and its contribution to her argument. Folding a piece of paper in half, for example, changes its shape. That change involves a physical rearrangement that puts pairs of edges of the

²We reserve the term *demonstration* for practical actions, though Clark (1996) uses the term more broadly to cover a wide range of communicative actions, including also iconic gestures and mimicry of quoted speech.

paper into spatial alignment with each other. The mathematics of origami abstracts the correspondences between edges paired this way as mathematical equalities between line segments.

Although Hart builds on our everyday understanding of demonstration, we argue in Section 2.3 that she has to do something special to endow those demonstrations with mathematical content. One problem is simply how different Hart's demonstrations are from the communicative actions we usually consider in semantics. We normally approach semantics by giving abstract rules relating symbolic expressions to meanings—this fits not only symbolic formal languages but arguably also speakers' grammatical knowledge of their native language (Montague 1970) and traditional systems of diagrams as well (Shin 1995). It's not clear how to decompose demonstrations such as Hart's folds into expressions of this kind. A more serious problem is the disconnect between the effects of demonstrations and meanings as traditionally conceived. We will endorse a received view, where our general understanding of events in the world must be qualitatively different from the kind of propositional content we appeal to in our arguments, especially in mathematics. Our understanding of events is incomplete, open-ended, and idiosyncratic. Meaning, by contrast, is regimented, precise, and public.

This paper explores a philosophical approach to demonstration that promises to defuse these puzzles. The approach is inspired by David Lewis's accounts of coordination and scorekeeping in conversation (Lewis 1969, 1979). As we see in Section 3.1, *coordination problems* are situations where agents share common interests, where agents' actions have to agree with one another, but where more than one successful strategy is possible. Coordination problems require agents to develop special relationships with one another, so they can establish aligned expectations and work together to make matching choices.

Communication naturally involves coordination. In principle, any action could be a signal of any meaning, as long as the agent and her audience expect the signal to be used that way. Coordination can just as well overlay practical actions with a distinctive further role in agents' collaborative relationships. Something like this is at play in Hart's origami mathematics, we suggest.

In particular, formal reasoning involves a particular kind of coordination: coordination on a *conversational scoreboard* (Lewis 1979) or *conversational record* (Thomason 1990).³ This is an abstract record of the symbolic information that interlocutors need to track in conversation. Among other things, the conversational record tracks what propositions interlocutors are committed to, how these propositions are related to one another, and what remains to be established to meet the aims of the

³Lewis developed his theory of the conversational record in a totally separate project from his work on coordination and convention. The two strands of research address rather different problems in the philosophy of language, and, perhaps surprisingly, Lewis himself never explored the connection between the two proposals. We believe, however, that there is a natural affinity between them. For example, combining the ideas makes it possible to characterize the domain of semantics in terms of conventions for updating the conversational record (Stojnic and Lepore 2014) and to show how such conventions provide standards for interpreting utterances in context (Lepore and Stone 2014). The development in this paper represents another combined application of Lewis's ideas.

conversation. The idea that we will explore is that Hart uses paper folding as an ingenious mnemonic to update the conversational record: each fold contributes specific mathematical inferences that abstract away from its concrete effects in reconfiguring the paper. To establish this conventional meaning for her actions, even with naive viewers, Hart draws on the general conventions of *discourse coherence*, which allow her to signal the correct interpretation of her actions by narrating what she does in suitable terms. Section 3.2 presents this framework in more detail, and explores the distinctive ways in which interlocutors combine words and actions when they coordinate situated conversation.

Section 4 revisits Hart's proof in light of these ideas. Our approach helps clarify the mechanisms Hart's proof appeals to, while revealing a range of problems for future research.

We focus on Hart's video because we think its mathematical precision (and elegance) helps to highlight the general representational function of demonstration in conversation. People regularly communicate their dynamic understanding of action in space and time by doing things. To get clear on what's happening in such cases, we think it's crucial to recognize that the critical actions have dual functions in fulfilling practical and communicative goals. The communicative effects are always distinctive, however, because only the communicative effects arise by linking up with interlocutors' mutual expectations about the conversational record. In demonstration, interlocutors also apply the distinctive knowledge they always need to communicate, but in a way that simultaneously depends on their practical knowledge of the common-sense world. This makes the phenomenon of demonstration a particularly clear and instructive example showing how communication draws not only on our knowledge of meaning but on our general capacities for planning, understanding and relating to others.


Of course, in many cases, speakers use demonstrations in ways that not only communicate precise ideas, but do so particularly effectively. For example, with demonstrations, speakers can prompt new insights and make their ideas particularly clear by linking up with their audience's perceptual abilities and their ongoing engagement with the physical environment. We conclude in Section 5 by emphasizing that this is an important factor of Hart's demonstration—one that remains to be explored in a full account of the power of demonstrations.

2 Hart's Proof

We begin with a tour of some mathematical ideas. We start with the Pythagorean theorem itself, which describes right triangles and establishes a relationship among the lengths of the three sides: namely, the square of the longer one is equal to the sum of the squares of the two shorter ones. Traditionally, we are given the lengths of the three sides as numbers a , b and c , in increasing order, and this relationship can be simply written $a^2 + b^2 = c^2$.

The result can be demonstrated formally using a variety of geometric and algebraic arguments. Hart's proof revisits a longstanding geometric construction that may in fact be due to Pythagoras himself (Maor 2007). The key idea is embodied in the

diagram in (2), which we construct by rotating and juxtaposing four copies of the target right triangle.



This construction turns out to create one large upright square with a small tilted square set precisely inside it. The four triangles and the larger square around them provide a framework to visualize the quantities the theorem relates. The tilted square has area c^2 . But it's obtained by removing four copies of the target triangle from the upright square. If we now raise rectangles framing two adjacent triangles, we get another way to visualize this area.⁴ But the remainder of the large square is now organized as an area a^2 and an area b^2 . Pythagoras himself seems to have demonstrated this result intuitively in terms of the physical rearrangement of four identical triangular tiles. Hart's proof uses origami to display the key geometric areas as regions of folded paper.


2.1 Diagramming the Pythagorean Theorem

Any proof in geometry ultimately involves using abstract operations to introduce mathematical objects in suitable relationships, as a way of witnessing some general truth. We're most familiar with Euclid's geometry, which allows the construction of diagrams such as (2) from primitive operations abstracted from the use of a straight edge and compass in drawing. Euclid's geometry is particularly easy to think about because its inferences are so naturally summarized as transformations in diagrams. The straight edge creates the line between two constructed points. The compass creates a circle centered at one constructed point and passing through another. Mathematically, invoking any of these operations constructs further points that become available from then on. Of course, we can undertake corresponding steps of drawing

⁴John Blackburne has published a simple animation presenting the insight of this proof at http://en.wikipedia.org/wiki/File:Pythag_anim.gif.

using a physical straight edge or a compass; images of the newly-constructed points appear at intersections in what we have drawn.

Consider the construction of (2), for example. We can realize this diagram by starting from an arbitrary right triangle, and proceeding three times through the copying process illustrated in (3).



The construction begins by extending the edges of the triangle. We use the compass to construct landmark points so we can raise a perpendicular copy of the hypotenuse of the triangle. With a different set of landmark points, we drop the final side of the triangle perpendicularly down.


In many respects, language provides an intuitive model for thinking about such diagrams, their content, and their role in mathematical argument. In particular, according to this received view, the diagram is an abstract type and the printed marks on the page are tokens of that type, just as an utterance is a token of the corresponding sentence.

In diagrams as in language, the type is constituted by the abstract elements involved and the formal relationships among them; it's these elements and relationships that determine meaning. With sentences, we are given the meanings to the words and combine the meanings compositionally based on their syntactic structure. To give the meaning of a diagram such as (2), meanwhile, we describe the abstract elements that make it up and the way they relate to one another. In (2), for example, we find sixteen abstract elements: eight line segments and eight intersection points. The formal relationships say which segments are involved in each intersection; as constructed there are further equalities among segments and among angles. These features are interpreted as specifying corresponding mathematical objects and relationships, just as the constructions call for. On this perspective, the difference between symbolic representations, including logical formulas or natural language sentences, and iconic representations, including diagrams, has to do with *how* structure is linked to meaning. Symbolic representations have a uniquely-determined, hierarchical syntax that determines their compositional semantics, whereas iconic representations have flat and possibly underdetermined structures but can still be interpreted through general rules that assign consistent meanings (Shin 1995; Greenberg 2013).

In diagrams as in language, there are many features of individual tokens that are irrelevant to determining types. One and the same sentence may be written or spoken, uttered loudly and quickly or softly and slowly, and performed by one or another speaker. Just so, one and the same diagram may be drawn with lines of different

thickness, in different colors, or at different scales. These features aren't relevant to identifying the type that (2) instantiates, according to the diagrammatic system conventionally used to present Euclidean geometry. Accordingly, they aren't interpreted as contributing to meaning.

This programmatic analogy suggests that giving a semantics for diagrams will involve the same architecture involved in giving a semantics for formulas. This suggestion has the salutary advantage of straightforwardly accommodating diagrams that feature symbolic annotations. Such cases are frequent and useful—consider (4), for example.



The diagram in (4) enriches the abstract underlying organization in (2) with indications of additional mathematical relationships. The mathematical content is determined, as always, as a function of the abstract elements that appear and the syntactic relationships among them. For example, the juxtaposition of certain segments with the symbols a , b , and c specifies the lengths of those segments in terms of variable quantities. Labeling an angle with a rectangular corner indicates that it is a right angle. Labeling an angle with a ticked arc indicates the measure of the angle in terms of a variable quantity.

In the diagram in (4), the annotations mark the inner and outer quadrilaterals as rectangles, and highlight equalities of further lengths and angles in the figure. These relationships are in fact established by the construction in (3). In general, mathematical relationships can be developed by inference and recorded in the diagram in tandem. In fact, the annotation in (4) summarizes everything we need to know to infer the Pythagorean theorem—we have an outer square that differs from the inner square by the addition of four identical triangles.

The symbolic annotations of (4) allow us to complete the proof of the Pythagorean theorem in explicitly symbolic terms, if we choose to do so. Measuring along the


outside edges, we see that the outer square has area $(a + b)^2$, which multiplies out to $a^2 + b^2 + 2ab$. We can also compose the area of the outer square by adding the area of the inner square, c^2 , to the area of the four triangles, which we can calculate as $4 \cdot \left(\frac{1}{2}ab\right) = 2ab$. The two expressions for the area of the outer square must be equal, so we derive

$$5 \quad a^2 + b^2 + 2ab = c^2 + 2ab$$

and now simplify to the explicitly Pythagorean (6).

$$6 \quad a^2 + b^2 = c^2.$$

However, it is also possible to visualize the relevant areas directly in the figure, as in (7).



First we divide the outer square into a lower part of height a and an upper part of height b . We further segment these regions to find a square of area a^2 and two $a-b-c$ right triangles, and a square of area b^2 and two more $a-b-c$ right triangles. So both $a^2 + b^2$ and c^2 represent the area of the outer square with four identical triangles removed.

2.2 Origami as Geometric Construction

A mathematician will tell you that a proof is any fair way of convincing someone of something, but to understand Hart's proof we need to get clear on its mathematical content and its inferential relationships. In the formal development of origami mathematics, due to Huzita and Justin (Huzita 1989; Justin 1989), folds play the role of marks made with straight edge and compass. Each fold constructs geometric objects

such as points and lines out of existing ones. However, setting up the fold in different ways yields different geometric results. You can make this vivid with a piece of paper with two marked points. Make a fold in the paper that passes through both points. This fold marks the line between the two points. Abstractly, this has the same effect as constructing the line between two points with Euclid's straight edge. We'll call this case (O1) following Huzita (1989) and Lang (2010). Unfold the paper and now make a fold (perpendicular to the previous one) so that the points coincide one on top of the other. The abstract operation corresponding to this new fold constructs the perpendicular bisector of the segment given by (O1). This is a new operation, indicated by (O2) in Huzita's and Lang's catalogues. Hart's proof invokes two other operations, which we invite the interested reader to explore further. Operation (O3) constructs a line equidistant between two lines with a fold that aligns one *line* onto another *line*. Operation (O4) folds a line over onto itself so as to raise a perpendicular segment; the location of this segment can be adjusted to make sure it passes through a specified point. These operations correspond to complex constructions that combine operations with straight edge and compass in Euclid's system. The Huzita and Justin system has a total of seven different folding operations—some not replicable with straight edge and compass—and completely characterizes the geometry of a single fold (Alperin and Lang 2009). We get even more powerful constructions by abstracting the effects of multiple folds undertaken simultaneously—common practice for enthusiasts of the art of origami.

Mathematically, origami's similarity to straight-edge and compass constructions gives it an analogous mathematical structure. For example, in both cases, algebraic techniques can be used to establish metatheorems about what can and cannot be constructed (Huzita 1989; Justin 1989; Alperin 2000). Origami happens to be somewhat more expressive than the straight edge and compass, in fact!

These mathematical preliminaries allow us to understand the mathematics behind Hart's proof, as well as the differences between Hart's origami and diagrams like those in (2), (4) or (7). We'll track the proof in detail here, as it unfolds in Hart's video. Unfortunately, we have to resort to figures that depict the changing state of the paper. Remember that Hart herself has no such figures, just the paper itself. We'd love to fold a piece of paper and show you what's happening, but the printed page precludes this strategy.

Hart sets the stage wryly with (8).


- 8 First you need a square, which you can probably obtain from a rectangle if you ask nicely.

As she says this, she folds an ordinary piece of paper along a diagonal (by O3), aligns the short and long edges of the paper (by O4), and rips a rectangle off the long edge at the level of the short one. Hart expects her audience to recognize the paper square that results from her activity; in fact, this square corresponds to the outer square of the diagrams in (2–7). You might notice many other results of this activity, such as the narrow strip of paper you get in addi-

tion to the square, or the texture that distinguishes the one ripped edge of the square from the three straight edges, but these effects are not relevant to Hart's argument.

Next, in tandem with (9), Hart folds the square into an isosceles right triangle eight layers thick (using O3 three times), yielding the shape in (10).

- 9 Step one: fold your square in half one way, then the other way, then across the diagonal.




Next comes the step that implicitly fixes the dimensions of the right triangle at issue in the proof (a use of the O4 fold).

- 11 Step two: make a crease along this triangle parallel to the side of the triangle that has the edges of the paper.

You get something like (12).

12




You might notice that Hart makes this crease by folding the top corner of the triangle up and over the bottom, rather than folding the top backward and down behind it; it transpires that nothing depends on this. In fact, it doesn't even matter for the proof *where* along the vertical edge of the triangle you measure out this fold! In (12), we've made the fold where it needs to be so we'll eventually wind up with an analogue of the diagram in (2). Hart not only remarks on the flexibility of this step, but dramatizes it in her demonstrations, by making this fold at two very different places in a first explanation of the proof and in a later summary.

At the next instruction (13), it's worth presenting Hart's narration in full.

- 13 Now when you unwrap it you'll have a square centered in your square. Extend those creases and make them sharp. And now we've got four lines all the same

distance from the edges, which will allow us to make a bunch of right triangles that are all exactly the same.


14



After unfolding and refolding the paper as Hart describes, what you get looks something like (14). Among the indefinite effects of this step, you can perhaps still see the traces of the folds you made and undid at the first step, dividing the square into eight pie wedges; these folds are shown in (14) as dotted lines. But your focus now is the four new creases you've just made, which you'll be working with for the rest of the proof. Perhaps, though, you might not have noticed the geometrical relationship among these creases, which will underpin the subsequent mathematics. In the formalism of the diagram (4), each line is at a distance a from the closer edge and a distance b from the further one. In describing (13), Hart remarks on the symmetry.


Hart now instructs us to fold back four outer triangles, matching the triangles visible in (2). We find them by folding diagonally across each corner from the crease closest to it to the crease closest to the next corner. (These are our first O1 folds, along lines between pairs of points.) We get a paper square with sides the length of the hypotenuse of those triangles, as depicted on the left in (15).

15




We juxtapose Hart’s folded square with a copy of the diagram in (4) to show that Hart’s folding now precisely displays the inner square of area c^2 presented in the diagram. As the depiction in (15) suggests with the dashed triangles, in the origami construction, we perhaps still have a sense of the triangles we have folded underneath and hidden from view.

Finally, Hart instructs us to fold back four triangles in a different pattern. Ripping a segment of the original lower right crease separates a small square tab. You can fold back the adjacent rectangle (two triangles) along one of the lines you’ve already constructed. You then fold back the rectangle above it and to the right (two more triangles, already constructed). In short, in this step, you return to the starting place of (14), as shown on the right in (16), then fold to create the two connected squares shown on the right in (16).



Compare Hart’s folded paper with the corresponding Euclidean diagram as elaborated in (17).



Hart’s paper displays two squares that correspond precisely to the a^2 and b^2 areas diagrammed in (17).

Just as with (7), the new folded shape in (16) makes clear that the square area defined by the hypotenuse occupies the same area as the two squares defined by the smaller sides put together. In fact, Hart proceeds to quickly alternate between the folded shapes in (15) and (16), highlighting the four equal triangles that are folded back in both cases and the equal areas that must therefore result. This equality is what Hart has set out to show.

2.3 The Puzzle

This abstract parallel suggests that the folded paper in Hart's video may have more in common with a Euclidean diagram than one might initially suspect. Using origami mathematics, we can think of a folded piece of paper, like a diagram, as an instance of an abstract type, constituted by formal relationships among its elements. In particular, in origami mathematics, we are to identify the concrete folds as instantiating corresponding abstract operations and interpret them through specific mathematical objects and relationships. In fact, you may decide that what Hart has done is to produce a diagram by folding it rather than by drawing it. The parallel certainly invites us to broaden our terminology this way—though we shouldn't underestimate the challenge of the dynamics of origami mathematics. Folding in origami mathematics is not a monotonic operation like drawing. Some of the folds of Hart's proof become possible only after previous folds have been undone.

Even if we think of Hart's proof as creating a kind of diagram, the creativity and physicality of her demonstration forces us to confront some of the problematic ingredients of the received story. How do viewers know that Hart's folded paper has this special interpretation? What establishes its mathematical meaning? More generally, how is it that any practical action in the world, of the sort that Hart exhibits, could figure in the symbolic economy we normally associate with language and those other actions we seem to use purely to express our thoughts?

The difficulty with demonstrations is the link between their content and their intrinsic effects. Hart seems to give her folds the content they have—creating lines of certain kinds—because of what the folds do—make creases with particular configurations in particular pieces of paper. But theorists across cognitive science have contested the idea that our understanding of practical action in the world results in anything like the propositional content associated with linguistic utterances and other communicative actions.

In the philosophy of interpretation, Dennett (1978) and Fodor (1983) call this the *frame problem*. Whenever you observe action in the world, there are an indefinite number of changes taking place, which you can only track using open-ended reasoning. We've seen this open-ended reasoning in the many different kinds of visible effects Hart's folds have, effects that we've remarked on in retracing her argument. If we can't list all the effects of an action, even in principle, then there can't be one right answer about what interlocutors should notice when an action takes place either. If demonstrations do have precise content, it can't just follow from the physical effects of the demonstrations alone; there must be other cues that determine the content of the speaker's argument. These cues must help viewers track the features of

the demonstration that matter, and help viewers not be distracted by features of the demonstration that don't.

In the psychology of interpretation, Sperber and Wilson (1986) develop a similar perspective. Confronted with a changing world, Sperber and Wilson suggest, each person focuses on the inferences that are easy for them to draw and that matter for them. In their theoretical terms, people attend only to *relevant* inferences, and those relevant inferences will vary from person to person and moment to moment. Actions will naturally carry different information for different people, and so—without some special cues—there won't normally be a particular propositional interpretation that you can associate with those actions.

Linguists will be familiar with the puzzle by way of Grice (1957), and his influential characterization of what's distinctive about linguistic meaning. For Grice, utterances appeal to intention recognition; the audience specifically aims to reconstruct the ideas the speaker has in mind. What audiences learn from observing an ordinary practical action, by contrast, seems to be just the information that's naturally carried by causality in the world. Grice's views again suggest that there's a fundamental divide between the "meaning" of practical action and the meaning of communicative actions. It's not clear how speakers can bridge this gap with their demonstrations.

In short, we need to think about demonstrations more carefully, in the context of the philosophy of communication, in order to explain exactly how practical effects can come to have the status of other propositional content that speakers contribute to the conversation.

3 Coordination and Coherence

We believe that the conceptual framework of Lewis (1969, 1979) offers a particularly useful set of tools to account for demonstrations as expressing precise meanings on a par with other communicative actions. Lewis's framework offers a general way of associating content with any action in the world, as long as interlocutors coordinate those actions, based on the right background of mutual expectations, and use them with the right dynamics of interaction, to update the conversational context. We review Lewis's ideas on coordination in Section 3.1 and his ideas on the nature of updates to conversation in Section 3.2. We apply these ideas to demonstrations in discourse in Section 3.3. Space precludes a detailed survey of the relationship of Lewis's ideas to alternative frameworks or the critical reaction to them. We refer the interested reader to Lepore and Stone (2014).

3.1 Coordination and Communication

Philosophers have long sought to naturalize meaning by relating it to a general capacity for social understanding that's applicable to all actions, regardless of their practical effects. According to Grice (1957), for example, we infer that an action has meaning when we recognize that its author has performed it with the intention of expressing an idea. (Here we gloss over some contentious technicalities

that will not concern us here.) Such accounts will naturally accommodate demonstrations, provided we can be precise about how viewers recognize the expressive intentions of demonstrations, and explain why this intention recognition underwrites the attribution of meaning.

In general, intention recognition has the same open-ended character that infects our understanding of practical action. We can never be sure we've completely grasped the dynamics of cause and effect that the agent has committed to in making the choices we have observed. We do sometimes have substantial information about how an agent has acted, we have good evidence of the agent's goals, and we can assume that the agent shares with us a realistic understanding of the relevant actions and their practical effects. In such cases, we can use our information to explain an agent's choices without any further reference to the agent's social context. For example, if we spot a colleague's car in front of us at the entrance to campus one morning, we can explain their subsequent choice of route by recognizing their intention to arrive expeditiously at their place of work. However, there is no limit on the factors that an agent may consider in choosing actions; intentions can be correspondingly difficult to recognize. Think for example of the habits, mnemonics and superstitions that could well be involved in explaining the specific parking spot your colleague takes.

To communicate, we must express our ideas in special ways—ways that wall off this open-ended inference. Lewis (1969) provides conceptual tools to explain how communicative intentions work and why they wall off open-ended inference. The key reason, Lewis argues, is that all coordination, including communication, involves a distinctive kind of inference and evidence.

Recall from Section 1 that coordination arises when agents' interests are aligned, where good outcomes require agents' actions to match, but where more than one successful strategy is available. Coordination problems are ubiquitous in the social world. For example, drivers face a coordination problem in sharing the road. They have a common interest in avoiding a crash. And they have multiple options for doing so, including driving on the right, as in the U.S., and driving on the left, as in the U.K.

The example illustrates why coordination is special. While we can anticipate others' routes just by knowing where they want to go, no such logic explains what side of the road they drive on. In general, we can't just look at agents' preferences and their beliefs about the practical effects of their actions and predict how they will act in social situations. Successful coordination depends on anticipating an apparently arbitrary choice that others make. Sometimes, of course, agreement could result from a lucky guess, or a quirk of biology, but cases like driving that are learned and variable cannot be explained this way. Rather, agents know the social context: agents rely on a shared background of mutual expectations to anticipate and understand one another. As we shall see, Lewis (1969) ultimately argues, following up ideas from Locke and Hume, that these mutual expectations are at the heart of the common-sense notion of convention.

For Lewis (1969), utterances are another case of action in the world that's taken under circumstances where agents have to coordinate. A key case is the signaling game—a situation where one party produces some action knowing the state of the world and the other party acts having seen the signal. Lewis's take on the Paul Revere story is an example of coordination by signaling. The sexton of the Old North Church

agrees in advance with Paul Revere that he will hang one lantern in the church tower if the British are coming by land and two if they are coming by sea. Revere will then ride off and organize the appropriate defenses. In the event, the sexton hangs one lantern in the tower of Old North Church. Revere rides off knowing that the British are coming by land and that he must defend against this approach. The pair could easily have adopted a different convention. For example, they could have used two lanterns to indicate that the British were coming by land. Had their strategies been different, Revere and the sexton would have acted differently—but they would still have accomplished their goals. That shows that they must coordinate with one another, and thus that they need a special kind of knowledge to decide how to act together. Of course, in this case, Revere and the sexton agree in advance on what they will do and how it is to be interpreted. This, for Lewis, is a prototypical case of a *convention*, but he characterizes convention more generally as a shared set of mutual expectations that allow a group to solve a coordination problem reliably. In the case of signals, the conventions are a repository for special interpretive knowledge that links them to their meanings.

In his formal development, Lewis defines signaling in a way that excludes demonstrations. Signals, for Lewis, can be produced with minimal costs and have no significant intrinsic effects on advancing agents' goals. By contrast, it matters what demonstrations do. It turns out, however, that this fact about demonstrations does not exclude them from Lewis's framework for communication and convention. The framework elucidates a role for coordination in actions that reveal the author's information and expectations, even when those actions simultaneously advance agents' practical interests.

Let's start with a purely practical case, which, like the case of driving on the right side of the road (or on the left), illustrates the ubiquitous need for agents to coordinate to achieve their goals in social situations. Two friends approach a closed door. They must pass through it one at a time; it will be awkward if each waits for the other or if both try to pass through immediately. They can coordinate this in different ways. The first to arrive can open the door, hold it while the second walks through, then follow afterwards. Or the first can open the door and walk through it, pushing back on the door so it stays open for the second to take it and go through behind.

Suppose the friends share a convention to hold the door rather than pass it. One friend arrives and opens the door. This naturally furthers the agents' practical interests, since someone must open the door if they are to pass through. But, in virtue of the friends' conventions, the first undertakes this action with a further intention, to hold open the door so the second can walk through. The same conventions allow the second to recognize this intention, and to act accordingly.

Conventions in this case play the same role in restricting the reasoning agents must do to recognize one another's intentions as they have played in our other examples. In particular, opening the door could play a role in an indefinite range of possible further plans. Opening the door sets one person up to hold the door for another and follow behind. Opening the door also sets one person up to walk through and pass the door to another. We can see this as an illustration of Dennett's and Fodor's frame problem: there are indefinitely many ways to describe what is happening when someone opens

a door. The possibilities reflect the genuinely ramifying practical consequences of the opening action.

At the same time, the example shows how agents can privilege certain of these practical consequences in their interactions with each other. Coordinating agents can recognize the intentions behind others' actions by appeal to convention, just as they might arrive at the right interpretation of a signal by appeal to convention. In both cases, agents try to understand one another by reasoning from mutual expectations that resolve ambiguities in how they will interact with one another. For our friends in this case, opening the door is part of holding the door, not part of passing the door. Despite the role for convention in intention recognition, in the case of practical actions, agents' coordinated strategies piggyback off actions' natural effects. The friends can only use opening the door as they do because this action has the practical results it has.

Focusing on coordination thus starts to defuse the puzzle of demonstration. Practical actions, like Hart's folding, *are* ambiguous: we can describe them in different ways, and so hypothesize different intentions on the part of their agents. But it's precisely for this reason that *using them meaningfully* is analogous to using utterances meaningfully. In both cases, we need a special kind of knowledge to explain how coordination succeeds. Coordination depends on the mutual expectations of the agent and her audience. We will suggest that demonstrations are like signals in that they are used to coordinate on information that furthers interlocutors' joint activity, and like practical action in that their contributions involve their practical effects. Before we can explain our view fully, however, we need to refine our understanding of how communicative actions contribute to conversation.

3.2 The Conversational Record

By linking communication to coordination, Lewis makes it easy to establish an interpretive parallel between utterances and demonstrations. But this isn't the whole story. Agents' coordinated strategies for action in the world, of the sort we appealed to in the case studies of Section 3.1, aren't fine-grained enough to zero in on the propositional content we ultimately want to associate with our communicative actions. Take the sexton's signal for example: does it mean that the British are coming by land? Or does it mean that Revere should ready land defenses? English distinguishes those meanings. Revere's and the sexton's strategies do not.

By contrast, when you look at conversation, it seems that people are coordinating not just on signals or information about what to do, but on the propositional content that they are contributing. We see this deeper coordination in the strategies interlocutors use in conversation, as they follow up their utterances with questions and clarifications, challenges and defenses. For example, interlocutors' inquiry into whether the British are taking a specific route will proceed in one way; their inquiry into how to defend against the attack will proceed in another way. This process can easily distinguish between fine-grained meanings. Put differently, it's only because the plan that Revere and the sexton have agreed to forecloses this inquiry that the status of its signals is indeterminate between these contents. To explain the meaning of utterances, diagrams, or demonstrations, in terms of Lewis's ideas of coordination

and convention, we need to characterize the process of coordinating on meaning in conversation in more detail.

Lewis's 1979 work on scorekeeping in conversation has played a seminal role in the development of our understanding of this process. For Lewis, interactions in conversations are governed by rules analogous to the rules of a game like chess or baseball. These rules describe what contributions interlocutors can make, how those contributions will be interpreted in context, and how those contributions will advance the context in turn. Of course, nothing restricts these rules to describe only linguistic utterances; they might apply equally to any of the actions interlocutors take in conversation. In Section 3.3, we'll see that we can and must conceive of demonstrations as making such rule-governed contributions to conversation.

As in a game like chess or baseball, we can think of the rules of conversation as establishing abstract distinctions that constitute the state of play. Lewis (1979) calls this *the conversational scoreboard*, but we prefer the term *conversational record* from Thomason (1990). The conversational record tracks the work interlocutors do to make, sharpen and defend their claims in conversation, and to use these contributions to pursue joint activity in the world. It tracks things like what has been said so far, what the agreed direction is for the conversation, what the topic is, what is salient for reference, what issues are open, what standards of meaning are in play, what each person is committed to so far and what everyone has agreed to. Roberts (2012) provides an overview of the linguistic evidence surrounding the conversational record and develops one particularly detailed proposal about the contents of the conversational record and its interfaces with the organization of conversation on the one hand, and the rules of grammar on the other.

Though the conversational record has proved more broadly applicable than Lewis may have envisaged, we think one of the important virtues of Lewis's analogy is sometimes lost in researchers' bids to expand on the approach. We should not expect the conversational record to track everything that happens in a conversation, or even everything that interlocutors recognize and share, any more than we should expect the scoreboard of a baseball game to track everything that happens on the field or everything that the players know. To make appeals to the conversational record as powerful as possible, our approach is to assume that the conversational record tracks just that information that is given a special status by the rules of language. The conversational record helps us to think precisely about the game that interlocutors play in carrying out a conversation, but we don't expect that this is all we need to say to describe the choices of interlocutors or the ways conversations tend to evolve.

The role of the conversational record is to track interlocutors' progress in specifying and agreeing on content. Thus, when we describe the dynamics of the conversational record, we fill in the details we need in order to link agents' actions to propositional content. For example, when one speaker makes an assertion, we record that the speaker is committed to the proposition. Others can then follow up by asking a clarification question, saying whether they agree or disagree, giving arguments for or against and so forth.

In order to fully develop a parallel between demonstration and other communicative actions, we need to go into some detail about how speakers—particularly speakers of English—actually do coordinate their updates to the conversational

record. We describe utterances as making abstract moves that effect specific changes to the conversational record. We call these moves *discourse relations*, or, more generally, *coherence relations* (Hobbs 1979, 1990; Kehler 2001; Asher and Lascarides 2003). Coherence relations provide general schemes for establishing implicit connections between utterances, and triggering corresponding changes in the evolving state of the conversation. For our purposes in this paper, a particularly revealing move is that of *Specification*, where a speaker follows up one utterance with a more detailed one, as a way of indicating exactly how the first claim is to be taken. *Specification* is a species of contiguity or elaboration relationship in the taxonomies of Kehler (2001) or Asher and Lascarides (2003). Example (18) gives an example of this move:

- 18 a. Chris and Sandy tidied the living room.
 b. They put away the toys, recycled the newspapers and vacuumed the carpet.

We imagine that a speaker offers (18a) in the course of an ongoing narrative, perhaps about the preparations for a party. The speaker continues with (18b). Normally, you'd expect a contribution to a narrative to say what happened next, but we don't interpret (18b) as describing what Chris and Sandy did *after* they cleaned the living room. Instead, on its most natural interpretation, (18b) lays out what precisely Chris and Sandy did *as part of* cleaning the living room. For example, we infer that the toys, newspapers and carpet described in (18b) started off in the living room (and were not initially in a tidy state).

Because *Specification* introduces interpretive dependencies between contributions to conversation, interlocutors can use it to amplify on the intended sense of vague or ambiguous language. Take (19).

- 19 a. Fido is a big dog.
 b. He weighs 35 pounds.

Again, the most salient interpretation of (19b) is as a *Specification* of (19a). On this interpretation, (19b) sums up the sense in which Fido is appropriately described by the term "big"—(19b) commits the speaker that the size of a dog is measured most directly by weight, rather than say height or length, and that a weight of 35 pounds definitely makes a dog big for the purposes of the conversation. (See Barker (2002).)

These commitments show that (19a) and (19b) are not understood as two separate interesting observations on the general theme of Fido's constitution. The meanings of (19b) and (19a) must be linked so that the information in (19b) backs up and underpins the speaker's claim in (19a). According to coherence theory, this connection is established by the speaker's use of (19b) in the relation of *Specification* to (19a). The rules of conversation recognize this as a particular kind of contribution to conversation, and interpret it by publicly registering the appropriate commitments on the conversational record. Thus, (19) shows how coherence can play a role in establishing fine-grained meanings for interlocutors' claims. In this sense, the example illustrates the value of coherence as a philosophical tool to avoid the interpretive underspecification otherwise inherent in coordination approaches to meaning.

Coherence not only enables speakers to resolve ambiguities and make vague language more precise, but it also enables speakers to associate new meanings with their communicative actions. This feature, we shall argue, is of particular relevance

in accounting for creative demonstrations like those in Hart's video. Example (20) offers an illustration from language use, inspired by the work of Clark (1983).

- 20 a. John did a Napoleon.
b. He put his right hand between the buttons of his coat and scowled.

Again, we naturally interpret (20b) as standing in a *Specification* relation to (20a). On this interpretation, (20b) provides information that spells out what's involved in the claim advanced in (20a). In effect, this tells us what the coinage "doing a Napoleon" must involve—affecting the idiosyncratic pose in which famous paintings depict the emperor.

Our view is that it is the rules of language that dictate that utterances such as (18b) are interpreted with these coherence relations.⁵ It's part of the game, just as it's part of the game that baseball pitches with certain trajectories are strikes, or that chess moves to certain configurations put the opponent's king in check. Utterances have to have the right content to make these moves, of course. But content alone does not dictate what moves utterances make; there also seems to be an element of convention. One way to see this is to consider the discourse in (21).

- 21 a. Chris and Sandy put away the toys, recycled the newspapers and vacuumed the carpet.
b. They tidied the living room.

Example (21) *does* read naturally as a narrative, in a way (18b) does not, especially if we imagine the speaker continuing either discourse with a further description of Chris and Sandy's efforts ("They also mopped the kitchen..."). With (21) we imagine that dealing with the toys, the newspapers and the carpet is preliminary to some indeterminate further tidying of the living room. Of course the sentences in the two discourses are the same, as are the logically possible relationships between them. Nor is the *Specification* meaning intrinsically ruled out by a speaker's purposes in adopting the ordering in (21): we can restore the *Specification* meaning in (21) by adding an explicit connective to (21b) to indicate it, such as "In other words". Ultimately, it simply seems like a matter of the mutual expectations of English speakers that we undertake a move of *Specification* by giving the claim first and spelling out what's involved in it afterwards.

Thus, there is a deep sense in which (20) illustrates coordination at work in establishing the meaning of "do a Napoleon". In understanding (20), interpreters must start from their shared knowledge of the rules of conversation, realize that the coherence relation of *Specification* applies, and reason from there to triangulate on the meaning of the speaker's neologism. As always in coordination, mutual expectations are decisive. The mutual expectations here say how utterances are to be associated with moves in discourse. They are decisive in the sense that they drive interlocutors' process of interpretation—with different strategies for achieving coherence, interlocutors would associate utterances with meanings in different ways. As we turn to

⁵Not all coherence theorists would agree with us here; Lepore and Stone (2014) explores the dialectic in more detail.

the interpretation of creative demonstrations, we will find similar evidence that principles of coherence are also what allow audiences to coordinate on the precise content of their contributions to conversation.

3.3 Coordinating Demonstrations

We can now state our main proposal: demonstrations, like the other communicative actions we have considered, get specific content because speakers use them to coordinate updates to the conversational record. In particular, speakers associate demonstrations with a special interpretation with a special status for their interaction. These coordinated contributions are constrained, but not determined, by the practical effects of the actions. Speakers can, and often do, intend both the communicative effects and the practical ones. Of course, this aligns with the understanding of Hart's folds we developed in Section 2.

This way of thinking about meaning brings a natural commonality between language and the gestures and practical actions which so often accompany it. But it requires us to get clear on the mutual expectations that allow interlocutors to coordinate on the contributions of demonstrations in our various examples. In particular, we need to understand the *coherence relations* that link demonstrations to ongoing conversation.

We regard this as a deep challenge—one that will require a substantial amount of future research. In particular, theories of discourse coherence offer specific, empirically-motivated accounts that inventory the possible moves in discourse, describe structural constraints that limit what moves are possible when, and specify the range of changes moves bring to the conversation. All these ingredients of coherence must be revisited for demonstrations, and the theoretical revisions are likely to be substantial. Lascarides and Stone (2009a, b), for example, propose to account for the interpretation of coverbal gestures by appropriate coherence relations. The proposal involves introducing new coherence relations, new possibilities for discourse structure, and new kinds of updates based on new distinctions in the conversational record. We expect a full account of demonstrations to involve further innovations of this kind.

For now, we consider a core case that epitomizes the promise and the challenges of the approach. We argue that there must be a special coherence relation *Compliance* that associates updates to the conversational record with actions that meet established expectations in context. For example, consider (22).

22 Can I borrow the key to the printer room?

In response, the addressee hands over an unlabeled plastic card. The addressee's action contributes to the conversation by complying with the speaker's request. We offer a number of arguments to think that *Compliance*, like *Specification*, is a coherence relation with an analogous role to other coherence relations in allowing interlocutors to coordinate their contributions to conversation.

For one thing, a *Compliance* relation helps to explain the special status that the action gets in the interlocutors' relationship. Roberts (2012) proposes that the conversational record is coupled with interlocutors' practical activity in part via a *to-do list*.

The to-do list itemizes all of the actions that interlocutors are committed to doing. One function of requests, even indirect ones such as (22), is to update the to-do list by adding something to it. After (22), the to-do list is augmented to indicate that the speaker has asked the addressee to provide access to the printer room.

Conversely, when the addressee of (22) actually provides access, the action is understood to discharge this pending task. The action must therefore be associated with appropriate updates to the conversational record. If we use the coherence relation of *Compliance* to link the interpretation of the addressee's response to (22), we fit these updates into Lewis's broader framework.

In fact, although it makes sense to update the to-do list in tandem with the action in this case, it's not the only option. If interlocutors had different expectations, they might have to update the to-do list only in tandem with an explicit claim by the addressee that the action has been taken care of, or a ruling by the speaker that the outcome was satisfactory. This is further evidence of the coordination we expect on Lewis's account.

In the case of (22), as in the case of (20), we see further evidence for coherence in the process by which the speaker recognizes the intention behind the addressee's action. Just as we must realize that (20) involves *Specification* to recognize that sticking one's hand in one's jacket and scowling amounts to "doing a Napoleon", we must realize that (22) involves *Compliance* to recognize that the piece of plastic the addressee is offering must be a swipe card that gives access to the printer room.

Empirical data from Clark and Krych (2004) offers a different kind of support for interpreting actions in conversation via coherence relations. Their data suggests that inferring *Compliance* plays much the same role in establishing and clarifying meaning in conversation as does inferring other coherence relations, such as the *Specification* relation we considered in Section 3.2.

Clark and Krych (2004) had pairs of subjects assemble Lego models. One was a director, who had a prototype model; the other was a follower, who had a pool of the pieces to fit together. Subjects' success in working together depended on the mix of action, gesture and demonstration they deployed in interacting with one another. In (23) for example (their Figure 6, page 72), director Sam tells follower Ted how to put two pieces together.

23 Get an eight piece green and join the two so it's all symmetric.

Sam breaks the utterance up into installments; in tandem, Ted gets the pieces as they're asked for, shows the pieces so Sam can verify that they are the right ones, and poises the pieces so that Sam can confirm their placement. Ted's step-by-step demonstration of understanding means that what looks like a single sentence—over just 8 seconds of this conversation—really involves a tightly-timed sequence of interactive turns where interlocutors give information, demonstrate their understanding, and make progress on their practical goals.

Ted's strategy depends on the interlocutors' mutual expectations that Ted's actions will comply with Sam's instructions. It is these expectations that allow Sam to use Ted's actions to understand whether Ted has understood the instructions correctly. By exhibiting the block, Ted shows what kind of block he takes Sam's "eight piece green" to describe. By poisoning it, Ted shows the configuration he

understands for Sam's "all symmetric". These moves commit the interlocutors to specific consequences of what Sam has said, much like the specifications of Section 3.2. Interpreting these actions thus means recognizing the coherence relations between Ted's actions and Sam's instructions and updating the conversational record accordingly.

The importance of coherence is even more vivid when interlocutors' mutual understanding becomes problematic. Clark and Krych (2004) offer an instructive example in (24), from their Figure 8, p. 74.

24 ...but only covers the last two— not those two— not that two— but— yes those two.

Here Clark and Krych's subject Susan offers this instruction to follower Tess at one point in their Lego assembly. Each time Susan breaks off, it's because Tess has poised the block in a new configuration. Over the course of Susan's seven second utterance, Tess poises the block in four different configurations. Only the last is correct. Although these actions do not all match what Susan actually wants Tess to do, Susan must still recognize that Tess is using them to indicate a possible interpretation of Susan's instruction. This way Susan knows to flag the mismatches in Tess's poising of the blocks and develop a more precise articulation of her instruction that the two can both understand.

Compliance is just one of many coherence relations that link utterances to actions. Hart's video regularly features a borderline case of *Compliance*, as Hart herself carries out the actions she is instructing in her own, parallel origami construction. Another possibility is that the utterance describes what the audience is supposed to observe in an accompanying event. We call this relation *Summary* in Stojnic et al. (2013). Engle (2000) offers a nice illustration in (25) (her Table 13, page 44):

25 You can probably hear it.

During a long pause, the explainer slowly moves the key into the keyhole: three clearly separate clicks can be heard.

Engle catalogued the demonstrations used by speakers to explain the working of a lock. With (25) the speaker makes the mechanism of the lock vivid to the audience by calling attention to the sound of the pins sliding into place as they actually do so.

We have nothing like a complete catalogue of these relations at this stage. We don't even know whether and how the relations fit in to the general taxonomies of coherence relations that have been developed in theories of discourse coherence. We also doubtless need to enrich our accounts of discourse structure and conversational state to account for utterances such as (23–25).

Despite this important caveats, we think our arguments offer strong evidence to conceptualize the role of practical action in conversation, including demonstration, in terms of coherence relations. Most importantly, we have suggested, the approach dovetails with a philosophically satisfying account of meaning that explains how a special sort of knowledge—convention—serves to associate practical actions with precise contributions to conversation. Interpretive reasoning requires us to integrate this conventional knowledge—across modalities—to come up with an overarching consistent pattern of contributions to conversation.

This integrated reasoning is already well documented for the coherent joint interpretation of speech and coverbal gesture (McNeill 1992; Engle 2000; Bavelas and Chovil 2000; Kendon 2004). For example, consider how gesture and speech relate in (26), as described by (Lascarides and Stone 2009a, Ex 16a, p. 168).

26 If you see this larger ball as ten small balls like that...

The speaker gets down off his chair to match his interlocutor, Susan, who sits across from him on the floor. His right arm now extends out in front of him at shoulder height, with his fingers curled and his index finger touching his thumb, as though holding a pen (an ASL 1-flat hand shape). During the gesture, the hand sweeps along a horizontal line further to the right.

The speaker of (26) is narrating a hypothetical physics experiment. By using the deictic expression “like that”, he shows that the gesture is going to help flesh out the spatial content of his utterance. But this speaker’s action does more than just specify how the balls are arranged. It creates a virtual space for the balls, and positions them in it. And his hand is posed to grip the object in a way that mirrors what an experimenter must do to carry out the physics experiment that he is implicitly describing.

As expected on our coherence theory, exactly the same sort of reasoning seems to be involved when interpreters make sense of utterances and other communicative actions used in tandem with demonstrations. In all cases, interpretive inference starts from a description of the observed action that leaves open key information. In communicative action, observers need to resolve ambiguities, fix underspecified meanings, and put references into context. In demonstration, likewise, observers may need to settle what action is being performed (moving, poisoning or exhibiting in 23) and recognize what objects are affected (a card or a key in 22). To establish coherence, interpreters must make assumptions and fill in missing information on the assumption that the action relates to ongoing events in the interaction in conventionally-established ways. Coordination is crucial, because it’s the mutual expectations of interlocutors that say what kinds of coherence relations are possible and how coherent contributions update the state of the conversation.

4 Exploring Demonstration

In Section 3, we sketched some arguments why we might want to account for demonstrations in terms of coordination, the conversational record, and principles of coherence. Promising as the account may be, there are clearly lots of gaps. In this section, we return to Hart’s proof, to indicate the work involved in spelling out a coherence account of demonstration and to make the ideas of Section 3 more concrete. In so doing, we hope to substantiate the parallels between demonstrations and other communicative actions, as well as the continuity between Hart’s folding and more familiar kinds of mathematical argument.

Looking more closely at Hart’s argument, we can pinpoint the key role for coherence in sharpening and clarifying the steps of her construction. In particular, Hart explicitly guides the viewer to the right mathematical understanding of her actions

by relating those actions coherently to the accompanying utterances, and thereby signaling the aspects of her activity that carry mathematical significance.

We start with (27).

27 Fold your square in half one way, then the other way, then across the diagonal.

Recall that Hart utters (27) while folding her paper square into an isosceles right triangle eight layers thick. Here as in (22), we have to find a coherent interpretation of the juxtaposed actions and utterances. In particular, with (27), we understand that Hart is doing something and describing it at the same time, via a kind of *Compliance* relation. Our characterization of this relation is clearly provisional—we know little about what kinds of utterances actions can target via *Compliance*, who can perform them, and how the relation is cued by the form of the utterance or the timing and manner of the action. However these issues turn out, recognizing *Compliance* must be crucial to understanding Hart’s argument, because the relation of *Compliance* so deeply informs our interpretation of the language and the significance we give the demonstration.

Hart’s instruction is underspecified; she doesn’t explicitly indicate that the first two folds follow the axes of the paper; she doesn’t say how to distinguish a single diagonal to guide the third fold. Only by watching the accompanying action do you realize that “one way” fits a horizontal fold that matches the top and bottom edges of the square, “the other way” fits a vertical fold that matches the right and left edges of the square, and “the diagonal” refers specifically to the diagonal of the the original square, now folded over into four coincident segments.

At the same time, language gives important clues to the idealized mathematical relationships Hart takes the folded paper to embody. For example, Hart’s “in half” highlights the fact that the (O3) folds she appeals to in this step have the mathematical result of creating equal lengths of paper. Hart continues her description of this step with the utterance “be precise”, lest there be any doubt that half is to be understood with specific mathematical significance. One way to capture the organization of Hart’s discourse is that “be precise” stands in an *Elaboration* relation to (27) as part of an extended discourse segment describing this sequence of folds.

As it happens, *Elaboration* is a relation that theorists often appeal to in accounting for the interpretation of extended discourse. However, in appealing to it here, we incur a variety of theoretical complications that we have yet to address. Can coherence relations link actions to multiple utterances at once? Can they link actions to extended discourse segments? What happens then? Only additional empirical investigation can settle such issues.

Hart’s description also hints that some aspects of the construction do *not* matter for the proof. For example, the paper square is symmetrical, so there’s no need to insist that the first fold is horizontal and the second vertical; it suffices to instruct the audience, as Hart does, simply to fold one way and then the other. To the extent that Hart can showcase the physical changes her argument abstracts, and downplay what’s not relevant, she mitigates the frame problem her audience would otherwise face in understanding her actions.

Hart’s instructions continue to exploit coherence to highlight the mathematical interpretation of the folds she makes. Hart’s (28) is an example. It’s another case of an

utterance followed up by an action in a *Compliance* relation. In this case, Hart explicitly notes that the (O4) fold called for at this step results in a precise mathematical alignment.

28 Make a crease across this triangle parallel to the side of the triangle that has the edges of the paper.

This alignment underwrites the use of this crease later in the proof to measure out identical copies of the Pythagorean Triangle around the side of the unfolded square.

Hart's voice-over not only includes moves that describe ongoing action in mathematical terms, via variants of *Compliance*, but also moves that highlight mathematical properties instantiated in the current situation, via variants of *Summary*. Hart exploits this to help her audience keep track of key aspects of the organization of the figure. One case where Hart exploits this device is in the setup of the main diagram, in (29), repeated from part of (13).

29 ... now we've got four lines all the same distance from the edges...

Hart offers this move as a mathematical description that follows from an abstract understanding of the folds she has performed. The single fold she just made in carrying out (28), parallel to the edges of the paper, now unfolds into four creases all aligned with the edges of the paper and all the same distance away. The coherence relation here is akin to the *Summary* relation illustrated in (25), but it has a special content in mathematical argument. As Hart articulates it in (30) at the end of her video, it's explicitly part of the contribution of such utterances that the audience should be able to verify their content on the basis of the geometrical operations that have already been carried out.

30 Of course, mathematicians are rebels and never believe anything anyone tells them unless they can prove it for themselves. So be sure to not believe me when I tell you things like "This is a square!" Think of a few ways you could convince yourself that, no matter what the triangles on the outside look like, this will always be a square, and not some sort of rhombus or parallelogram or dolphin or something.

We see Hart's commentary as a self-conscious introduction to the background expectations of mathematical practice and discourse—precisely those principles that, we have argued here, give her the resources to coordinate with an engaged audience and demonstrate the idea she has in mind. Thus, though it's clear what Hart is doing with (29), regimenting it in theoretical terms is challenging. Is there a special variant of the *Summary* relation involved in mathematical demonstration? If so, what basis should we give for our taxonomy of coherence relations? If not, how do we account for the mathematical force that Hart summarizes in (30)?

Hart's use of presupposed meaning in (28) illustrates one final strategy for using coherence to establish mathematical content. The noun phrase "the side of the triangle that has the edges of the paper" encodes the presupposition that there is exactly one such side. Presuppositions represent background information that's expected to be readily accessible and uncontroversial to interlocutors. However, presuppositions can still call attention to information that the audience might not otherwise have

noticed. Though Hart takes it for granted that one side of the triangle has the edges of the paper, this might not have occurred to you just by watching the proof the first time. We might appeal to the mechanism of presupposition accommodation (Lewis 1979; Thomason 1990) here, and suppose that it is Hart's utterance in (28)—and not the folding that she does—that triggers interlocutors to add the fact that there is exactly one such square to the conversational record.

The video shows that Hart knows that this is a place where the proof can go wrong if you're not paying attention. While she says (28), she fans the edges of the paper so that you can see what she's talking about. We're not sure whether to think of this action as a demonstration, or just as the flashy performance of the kind of deictic gesture explored in Lascarides and Stone (2009b). It seems to depend on how you think about the effects of the action. Deciding whether and how to draw such distinctions represents yet another challenge for future research.

However we think of Hart's action here, her choice to use it suggests that the puzzle of demonstration that we opened with is a real one. The effects of action are indefinite—even in a highly-structured case like Hart's proof. To make sure that the effects of actions become a part of the public record of the conversation, interlocutors have to put expectations in place that privilege those effects. One way to do that is by announcing in advance what the action will do; in this case the action is interpreted coherently to link up with established expectations. (28) shows another way: signaling after the fact that some effect that's already been achieved, but perhaps not noticed, really matters for the conversation.

5 Conclusion

In face-to-face conversation, speakers use all the means at their disposal to get their ideas across. They talk, they gesture, but they also carry out practical actions in the world. Who could imagine a good explanation of a complex movement, of how to dance a step or strike a pose, of how to tie a knot, of how to throw a frisbee, of how to use a knife, that did not involve an accompanying performance of the actions themselves?

Previous formal and empirical analyses generally depend on classifying actions in conversation either as communicative actions or as instrumental—despite the obvious possibility for overlap. In contrast, we have pursued a perspective on meaning in conversation where we can describe a single action both in terms of the concrete, useful changes that it brings about in the world and the distinctive contributions it signals in the conversation. Speakers imbue demonstrations with meaning by exploiting the principles of coherence in discourse. Coherence relations regiment our conventional expectations about how to get ideas across to one another. We normally act to meet these expectations. So these expectations can give the actions that we take to meet them clear meanings—meanings that we can think of as being analogous to linguistic meanings even when the actions are not utterances, but the many other kinds of actions we take in the service of sharing ideas and getting things done.

We have contrasted the communicative force of demonstration with the open-ended engagement that is sometimes required to fully appreciate language and action

alike. An important part of the appeal of visual explanations is the visual thinking they recruit, and Hart's proof is no exception. We've already remarked on the power of folding the different areas as a perceptually vivid way to show the relationship between them. As Hart repeatedly flips back and forth between the c^2 (15) and a^2+b^2 (16) configurations of the paper, she invites her viewers to *see* that these areas are the same.

But Hart also recruits her viewers' detailed visual engagement moment by moment through her explanation. Take (31), for example, another part of the key instruction in (31).

31 Now, when you unwrap it, you'll have a square centered in your square.

Hart's (31) is another summary of the mathematics of the figure. But it also invites the viewer to do some visual thinking. In understanding (31), the viewer can draw on the mechanisms of perceptual grouping to reanalyze the unfolded paper into a new pattern. If you succeed in seeing this pattern, you have a new and intuitive way to keep track of the complex mathematical relationships that bind the proof together. Though we've focused here on the fact that demonstrations can convey mathematical content, that should not be taken to diminish the many other ways that seeing something happen can prompt inference and insight.

Acknowledgments This research was supported in part by NSF IIS-1017811. Preliminary versions have been presented at colloquia and UCLA, CUNY, Edinburgh and Rutgers, and in talks at the Amsterdam Colloquium–Semdial joint session and the Rutgers–Jagiellonian Conference on Cognitive Science. This presentation benefits from comments from audiences there, from Sam Cumming, Doug DeCarlo, Eileen Kowler and Rochel Gelman, and the referees of this issue, and particularly from extensive discussion with Gabe Greenberg and Ernie Lepore.

References

- Alperin, R.C. 2000. A mathematical theory of origami constructions and numbers. *New York Journal of Mathematics* 6: 119–133.
- Alperin, R.C., and R.J. Lang. 2009. One, two and multi-fold origami axioms. In *Origami*, ed. R.J. Lang, 4, 371–394. A K Peters.
- Asher, N., and A. Lascarides. 2003. *Logics of conversation*. Cambridge: Cambridge University Press.
- Barker, C. 2002. The dynamics of vagueness. *Linguistics and Philosophy* 25(1): 1–36.
- Bavelas, J.B., and N. Chovil. 2000. Visible acts of meaning: An integrated message model of language in face-to-face dialogue. *Journal of Language and Social Psychology* 19(2): 163–194.
- Clark, H.H. 1983. Making sense of nonce sense. In *The process of language understanding*, eds. G.F. d'Arcais, and R. Jarvella, 297–331. London: John Wiley and Sons.
- Clark, H.H. 1996. *Using language*. Cambridge.
- Clark, H.H., and M.A. Krych. 2004. Speaking while monitoring addressees for understanding. *Journal of Memory and Language* 50: 62–81.
- Dennett, D.C. 1978. *Brainstorms*. Bradford Books.
- Engle, R.A. 2000. *Toward a theory of multimodal communication: Combining speech, gestures, diagrams and demonstrations in instructional explanations* Thèse de doctorat non publiée. Stanford University.
- Fodor, J.A. 1983. *The modularity of mind: An essay on faculty psychology*. MIT.
- Greenberg, G. 2013. Beyond resemblance. *The Philosophical Review* 122(2): 215–287.
- Grice, H.P. 1957. Meaning. *The Philosophical Review* 66(3): 377–388.
- Hobbs, J. 1990. *Literature and cognition*. CSLI.
- Hobbs, J.R. 1979. Coherence and coreference. *Cognitive Science* 3(1): 67–90.

- Huzita, H. 1989. Understanding geometry through origami axioms. In *Proceedings of the first international meeting of origami science and technology*, ed. H. Huzita, 215–222.
- Justin, J. 1989. Resolution par le pliage de l'équation du troisième degré et applications géométriques. In *Proceedings of the first international meeting of origami science and technology*, ed. H. Huzita, 251–261.
- Kehler, A. 2001. *Coherence, reference and the theory of grammar*. CSLI.
- Kendon, A. 2004. *Gesture: Visible action as utterance*. Cambridge.
- Lang, R.J. 2010. Origami and geometric constructions. Online Tutorial. <http://www.langorigami.com>.
- Lascarides, A., and M. Stone. 2009a. Discourse coherence and gesture interpretation. *Gesture* 9(2): 147–180.
- Lascarides, A., and M. Stone. 2009b. A formal semantic analysis of gesture. *Journal of Semantics* 26(3): 393–449.
- Lepore, E., and M. Stone. 2014. *Imagination and convention: Distinguishing grammar from inference in language*. Oxford.
- Lewis, D.K. 1969. *Convention: A philosophical study*. Harvard.
- Lewis, D.K. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8: 339–359.
- Maor, E. 2007. *The pythagorean theorem: A 4,000 year history*. Princeton University Press.
- McNeill, D. 1992. *Hand and mind: What gestures reveal about thought*. Chicago.
- Montague, R. 1970. Universal grammar. *Theoria* 36(3): 373–398.
- Roberts, C. 2012. Information structure: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics* 5(6): 1–69.
- Shin, S.J. 1995. *The logical status of diagrams*. Cambridge: Cambridge University Press.
- Sperber, D., and D. Wilson. 1986. *Relevance: Communication and cognition*. Harvard.
- Stojnic, U., and E. Lepore. 2014. What's what's said? In *What is said and what is not: The semanticspragmatics interface*, eds. C. Penco, and F. Domaneschi, 17–36. CSLI Publications.
- Stojnic, U., M. Stone, and E. Lepore. 2013. Deixis (even without pointing). *Philosophical Perspectives*. 26(1): 502–525.
- Thomason, R.H. 1990. Accommodation, meaning and implicature. In *Intentions in communication*, eds. R. Cohen, J. Morgan, and M.E. Pollack, 325–363. Cambridge: MIT Press.