# How do you Assert a Graph? Towards an account of depictions in scientific testimony

Corey Dethier

[forthcoming in *Noûs*]

**Abstract**

I extend the literature on norms of assertion to the ubiquitous use of graphs in scientific papers and presentations, which I term "graphical testimony." On my account, the testimonial presentation of a graph involves commitment to both (a) the in-context reliability of the graph's framing devices and (b) the perspective-relative accuracy of the graph's content. Despite apparent disagreements between my account and traditional accounts of assertion, the two are compatible and I argue that we should expect a similar pattern of commitments in a set of cases that extends beyond the graphical one. I end by demonstrating that the account resolves apparent tensions between the demands of honesty and the common scientific practice of presenting idealized or simplified graphs: these "distortions" can be honest so long as there's the right kind of alignment between the distortion and the background beliefs and values of the audience.

## 0 Introduction

In 1999, *Geophysical Research Letters* published a short paper (Mann, Bradley, & Hughes, 1999) on climate change in the Northern Hemisphere. Included was a graph that looks vaguely like a hockey stick: a long relatively flat period followed by a sharp uptick. Over the next decade, versions of the "hockey stick graph" became both one of the main tools for communicating the novelty and seriousness of human-induced climate change and one of the main targets of denialist attacks. When the dust settled in the late 2000s, Mann, Bradley, & Hughes (1999) had been vindicated: subsequent studies have consistently replicated the general hockey stick shape and the conclusion that we're currently facing the highest temperatures in thousands of years.

The furor over hockey stick-shaped curves in the 2000s was exceptional for the amount of political attention that it garnered—see Mann (2013)—but not for its focus on a graph. Graphs—and other depictions like figures, diagrams, maps, tables, charts, and pictures—are ubiquitous in the sciences; they're particularly important in the context of communicating complex scientific ideas to non-experts. And there are ways of abusing these tools—books on graphical best practices, such as Kosslyn (2006) and Tufte (2001), devote whole chapters to "lying" with graphs. Fiddle with the scales or axes in the right way, and your graph will appear to show the exact opposite of what the evidence actually supports. Given that depictions like graphs are both commonplace in scientific testimony and capable of rendering that testimony misleading, our philosophical accounts of scientific testimony should address their use.

To date, however, the philosophical literature on scientific testimony has largely focused on declarative sentences. Indeed, the philosophical literature on testimony in general has typically ignored testimony involving non-linguistic vehicles. While there's some work outside of epistemology that touches on the subject, there is little extended discussion of the relationship between depictive testimony and its more familiar linguistic cousin.[1] I aim to bridge this gap, at least for the case of graphs. My contention is that in cases of what I'll call "graphical testimony," the speaker's epistemic commitments extend to non-truth-apt features of the graph—such as the scale—that affect the *perspective* that the audience adopts towards the graph's contentful elements.[2] More specifically, the account that I defend is that a speaker who uses (as opposed to mentions) a graph is epistemically responsible for both content and their perspective-influencing choices: the latter must contribute to a reliable perspective and the former must be accurate from that perspective. Each of these commitments is of a kind with the commitment involved in assertion: if the knowledge norm provides the right account of the norms governing assertion, then a speaker should present a graph only if she knows both that her choices contribute to a reliable perspective and that the content is accurate given that perspective.

I begin by sketching a simple picture of the norms governing graphical testimony and raise two problems for it (§1). I then show that both problems can

[1] Though two exceptions—García-Carpintero (2023) and Lewerentz & Viebahn (2023)—have appeared since I began writing this paper. There are quite a few points of intersection between the account offered here and the work of Lewerentz and Viebahn in particular, but unfortunately considerations of space prevent a thorough comparison at present.

[2] I'm borrowing the technical terminology of perspectives from Elisabeth Camp. See Camp (2019) for an introduction and section 2 for details. While Camp is obviously drawing inspiration from similar sources as the literature on "perspectivism" in philosophy of science (see, e.g., Giere, 2006; Massimi, 2022; van Fraassen, 2008), neither her use nor mine should be understood in those terms.

be resolved by incorporating perspective (§2). In the second half of the paper, I argue that my account is compatible with traditional views on assertion and suggest that we should expect perspectives to play a similar role in a wider class of cases (§3). Finally, I end by demonstrating that the account resolves apparent tensions between the demands of honesty and the common scientific practice of presenting idealized or simplified graphs: these "distortions" can be honest so long as there's the right kind of alignment between the distortion and the background beliefs and values of the audience (§4).

# 1 Graphical testimony and assertion

"Assertion" is a term of art in philosophy: at minimum, to *assert* P is to say that P is true. Typically, philosophers include further conditions; the paradigmatic assertion is one where the speaker utters or writes a sentence that expresses P with the intention that their audience believe that P is true. So a sneeze might communicate that I am ill, but it is not an assertion because (a) the act of sneezing evidences rather than states my illness and (b) sneezes are (typically) unintentional rather than intended to communicate.

One of the primary philosophical projects surrounding assertion is determining the criteria, standards, or (as it is usually put) "norms" that we (should) use to evaluate assertions. Typically, the suggested norms take the form of a necessary condition: an assertion of P is proper (as an assertion) only if [...]. Alternative accounts fill in the blank in different ways: P is true (Weiner, 2005), the speaker believes that P is true (Bach, 2008), the speaker would be justified in believing that P is true (Lackey, 2008), the speaker knows that P is true (Williamson, 2000), etc. Common to each of these accounts is epistemic commitment: all of them view speakers as committed to (at least) the content of their assertions, and disagree (only) about the nature of that commitment.[3]

To date, most of the philosophical literature on scientific testimony has been concerned with the use of assertions in science.[4] For example: Dang & Bright (2021) and Fleisher (2021) argue that the declarative sentences found in (some parts of) scientific testimony are not governed by any of the norms canvased above; Buckwalter (2023) and Dethier (2022a) disagree. As Gerken (2022, 18) notes, however, scientific testimony is not exhausted by paradigm examples of

[3] A separate—though related—project concerns the nature of the norm or norms governing assertion; for an overview, see the supplement to Pagin & Marsili (2021). It's an open question whether graphical testimony has interesting implications for this second project.

[4] That's not quite right: there's significant discussion of the role of values in science communication that's not explicitly focused on assertion. I discuss one connection with these debates in section 4.

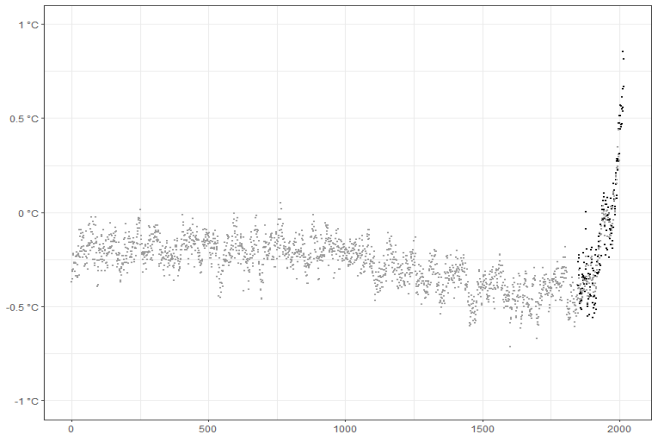


Figure 1: An example of a "hockey stick" graph showing deviations from 20th century average global temperature. Grey dots represent historical reconstructions, black dots observations. Generated using data from PAGES 2k Consortium (2019).

assertion. On the contrary, scientific testimony frequently involves depictions such as graphs, charts, diagrams, figures, pictures, tables, and maps. Our focus in this paper will be on graphs, though we'll see that there are reasons to expect that the lessons generalize. In any case, call testimony by way of graphical presentation "graphical testimony." The aim of this paper is to develop an account of the norms of graphical testimony akin to accounts of (linguistic) assertion.[5]

There's disagreement in the literature on linguistic testimony over whether a speaker's epistemic commitments extend beyond the content of their assertions to insinuations and implicatures—contrast, e.g., Fricker (2012) with Camp (2018a) or Viebahn (2020). There are some cases, however, that seem to fall clearly outside the scope of epistemic commitment. For example: while a speaker might be criticized for uttering a sentence in English when her interlocutors only speak German, her fault is pragmatic or ethical, not epistemic. Even insinuations are truth-apt, and thus open to epistemic commitment. By contrast, the choice of language is not truth-apt, and so is not open to criticism on the grounds that it isn't true or that the speaker doesn't bear the right epistemic attitude towards it.

These cases suggest that the distinction between those elements of a graph that are truth-apt and those that aren't provides a plausible first pass on the

[5]There are potentially gaps between the norms that govern testimony and those govern assertion; see Lackey (2011) for one way this kind of gap could arise. As indicated, our focus will be on those norms of graphical testimony that parallel the norms of assertion rather than on the (potentially) broader category.

bounds of epistemic commitment in graphical testimony. Consider figure 1. We can think of figure 1 as divided into two parts: the space of possible values—in this case, a Cartesian plane whose coordinates correspond to year-temperature pairs—and markings—in this case, dots—that distinguish the actual values from the merely possible ones. The plane is not truth-apt; it's more like a language than like a proposition. By contrast, the individual markings are more like propositions. A grey dot placed at coordinates $(1, -.369)$ can be thought of as expressing a proposition like [[the reconstructed mean global temperature for 1 CE is .369°C below the $20^{\text{th}}$ Century average]]. Our first pass, then, would yield an account on which someone who testifies by presenting this graph is epistemically committed to the accurate placement of each dot, but not to features like the scale of the graph. If the knowledge norm of assertion is the right account of the norms of testimony, then she must know the proposition that corresponds to each dot, but she need not know that the scale is true, because (after all) that's a category mistake.

Given that the sketchy account just offered is a first pass, we should expect that there are some details to be worked out. For instance, while hockey stick graphs are typically used to communicate the novelty of climate change (see IPCC, 2021, 6), figure 1 doesn't actually depict novelty: there's no sense in which the novelty of climate change could be incorrectly plotted. Instead, novelty is something that the audience sees for themselves when they look at the graph, and thus more like an insinuation or an implication than like the content of an uttered sentence.[6] The account so far is silent on whether the speaker is responsible for that claim. Similarly, there are questions about whether the "content" of the graph includes (e.g.) any "gaps" between dots and whether the speaker is responsible for the lack of a dot in a particular location. Questions like these are familiar from the literature on the semantics of depictions, and a full account of the epistemic responsibilities involved in graphical testimony almost certainly requires a more fleshed-out semantics for graphs than I've offered so far.[7]

Some of the problems with the account go beyond "details to be worked out," however. The first of these problems is straightforward. Like any other aspect of science, the depictions found in scientific testimony often involve substantial "distortions": simplifications, idealizations, abstractions, etc.[8] Consider figure 2,

[6]Notably, the IPCC spells the conclusion out explicitly in IPCC (2021, 6); there the graph itself serves more to illustrate the unprecedented character of climate change than to communicate it.

[7]Work on the semantics of depictions is necessarily focused on particular kinds of depictions, of which maps and pictures have received the most attention. For examples, see Camp (2018b) and Greenberg (2021) respectively. The literature on the semantics of graphs specifically is almost non-existent, though Perini (2005) is an exception and the literature on diagrammatic proof (e.g. Azzouni, 2013; de Toffoli, 2017) touches on graphs as well.

[8]The presence of distortions in the declarative sentences of scientific testimony is arguably

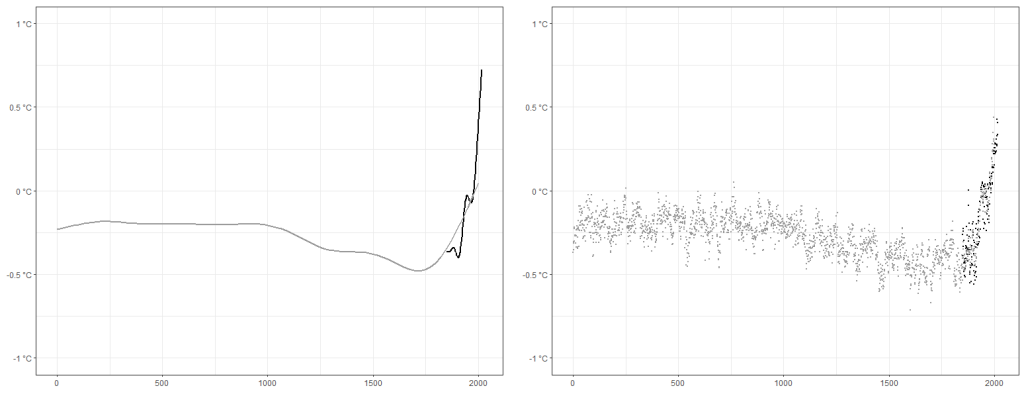


(a) Smoothed data (b) Manipulated data

Figure 2: Two distorted versions of figure 1. The first has been smoothed; the second has had approximately 60 "outliers" on the far right of the graph manipulated downwards to create the impression of less dramatic upswing.

which provides two distorted versions of figure 1. Applying the same principles to either of the graphs in figure 2 that we applied to figure 1 yields epistemic commitments to propositions that are strictly-speaking false: the line in figure 2a passes through coordinates that don't correspond to any real temperature-year pairings and approximately 60 of the roughly 2100 data points in figure 2b have been manipulated downwards.

Despite the fact that in an important sense the first graph is more distorted than the second, it's the one that could conceivably be used in epistemically permissible testimony. In some contexts, the smoothing methods employed in figure 2a would be treated as misleading—smoothing out market fluctuations might mislead the risk-averse investor, for example[9]—but they often aren't, and probably wouldn't be in this case. By contrast, the manipulations found in figure 2b render it ineligible for use in permissible testimony (excluding the case where the testimony concerns the presence of distortion); even though the effect is much milder and the second graph correctly plots much more than the first, systematic manipulation of data is intolerable in a way smoothing data isn't. So, as illustrated by figure 2a, graphical testimony can be epistemically permissible even when the truth-apt markings on the graph don't "line up" with true propositions. And we can't explain this permissibility by appealing to something like a contextually-lowered standard for presenting the graph, because figure 2b will meet any standard of accuracy that figure 2a does.

There are two potential diagnoses of this problem. On the one hand, it might

a problem for traditional accounts of assertion as well. See section 3 for discussion.

[9]I owe this example to Mikkel Gerken.

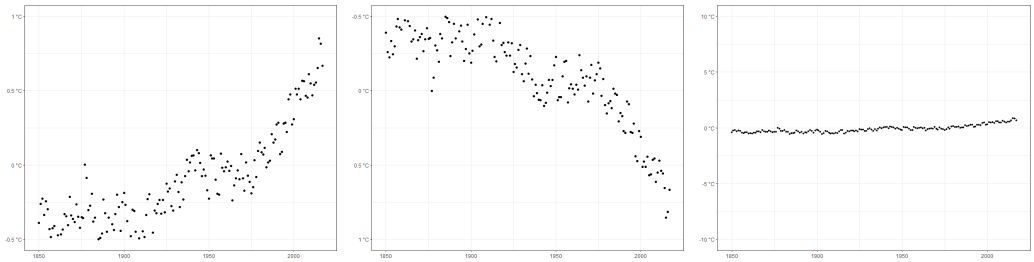

(a) "Standard" y axis (b) Reversed y axis (c) Expanded y axis

Figure 3: Three graphs of global temperature changes during the period 1850-2017. The first graph, (a), has a standard y axis. In (b), the y axis is reversed—it counts down rather than up. In (c), the scale on the axis has been expanded. Generated using instrumental data from PAGES 2k Consortium (2019).

be that our pre-theoretic assumptions about how to relate markings on the graph to propositions are too simplistic, and that a proper semantics for graphs would reveal that figure 2a doesn't actually "say" anything false. On the other hand, the problem might be a fault in our account of the norms governing graphical testimony: on this reading, good graphical testimony sometimes contains truth-apt elements that are strictly-speaking false. Regardless of which option we choose, however, the presence of distortions in permissible graphical testimony requires a substantial revision to the account offered above.

The second problem is that choices regarding non-truth-apt elements such as the scale on a graph have consequences for how the graph appears or is perceived. Consider the simple and unsubtle examples provided in figure 3. These three graphs are all constructed using the same data, and in all three cases the data points are placed accurately. Nevertheless, at least to the untrained eye, they present three very different pictures of climate change: the first graph shows us that temperatures are rapidly increasing, the second that they're decreasing, and the third that they're basically constant.

In my experience, at least, scientists are willing to level epistemic criticisms at choices like these, even though they operate on non-truth-apt elements of the graph. Here's the climate scientist Gavin Schmidt, for example, criticizing choices regarding the "baseline," which is essentially the x-intercept used when comparing different trend lines:

> Worrying about [the] baseline used for the anomalies can seem silly, since trends are insensitive to the baseline. However there are visual consequences to this choice. Given the internal variability of the system, baselines to short periods (a year or two or three) cause larger spreads away from the calibration period. Picking a period that was

> anomalously warm in the observations pushes those lines down relative to the models exaggerating the difference later in time. (Schmidt, 2016)

Schmidt's complaint, in essence, is that the choice of baseline—while strictly-speaking non-truth-apt—causes the graph to misrepresent important features of the target, which in this case is the difference (that is, the "spread") between two trend lines as time increases. For Schmidt, at least, a graph with a bad scale isn't akin to an impolitic but true assertion; on the contrary, he opens his post by describing the graphs he's criticizing as "misleading" and promises to offer a "more honest" presentation.

The scale is hardly the only aspect of the graph that the speaker can manipulate to alter how the graph appears. Constructing a graph involves dozens of decisions about non-truth-apt elements: we can choose to include more data or less, to use different types of graphs, to smooth the data, etc. All of these choices *can* affect how the graph is interpreted and read, and the speaker of a graph is appropriately treated as epistemically responsible for any of them that do. After all, a layperson without substantive background knowledge about climate change has to take the choice of scale on trust in much the same way that they have to take the depiction of the data. If the audience is to gain knowledge from the speaker's testimony, both elements must be chosen correctly.

To reiterate, the problem here is not that scientists sometimes criticize the choice of scale. It's that the distinction between truth-apt and non-truth-apt elements sketched earlier is permeable and, accordingly, scientists don't seem to draw a clear line between the (epistemic) responsibility that the speaker incurs in placing the data and the responsibility that she incurs in choosing the scale. In order to account for this feature of our practice, either the scale must somehow affect the content of the depiction despite not itself being truth-apt or the speaker's responsibility must go beyond said content. In either case, we need to amend the simple picture offered above.

## 2 Perspectives in graphical testimony

What we need are more resources: we need tools to distinguish both between permissible and impermissible distortions and between permissible and impermissible choices regarding features like scale. Fortunately, both needs can be met by concepts developed by Elisabeth Camp (2017, 2018a, 2019) in discussing the use of metaphors, insinuations, and other kinds of not-quite literal language.[10] As Camp points out, a statement like Romeo's famous "Juliet is the Sun" serves not

[10]My use of Camp's work is inspired by the similar appeals found in Fraser (2021).

as literal predication but instead communicates Romeo's (current) worldview—not just what he believes, but also what is currently important to him and how his experience is organized (Camp, 2017, 50). Camp calls this worldview a "perspective": "an open-ended disposition to characterize: to encounter, interpret, and respond to some parts of the world in certain ways" (Camp, 2019, 24). Graphs frequently come with or are encountered through a perspective; in figure 3b, for instance, the inclination to "see" decreasing temperatures is (part of) a perspective in Camp's sense.

"Frames"—or sometimes "framing devices"—are the representational vehicles that convey perspectives (Camp, 2019, 27). The metaphorical comparison of Juliet to the sun, for example, or the invocation of a stereotype. Much of what scientists do in designing a graph involves "framing" in this sense. In our running example, the scale (quite literally) provides a frame on the data and, as we just saw, manipulating this frame can alter the perspective that viewers adopt (or that is forced on them) when they encounter the graph. Similarly, choices about the degree of smoothing, which data to include and exclude, whether to indicate confidence intervals, even what colors to use—all of these can shape the perspective that the audience adopts and thus the conclusions that they take away.

Of course, something like the scale doesn't serve to "express" a perspective in the same way that "Juliet is the Sun" does: the scale doesn't have content of its own. But Camp is clear that she intends framing devices to go beyond traditional expressive sentences, and the graphical frame that's constituted by the choice of scale and all of the other choices that go into shaping how a graph looks affect how viewers "encounter" the content of the graph; they influence the perspective that viewers adopt when interacting with the depicted content. And my claim in this section is that a speaker who presents a graph is just as epistemically responsible for these framing choices as she is for the truth-apt content.[11]

I'm getting ahead of myself, however. The point of bringing in Camp's technical machinery is to resolve the problems with the simple account of graphical testimony outlined in the last section. Recall that the first problem was that simplifications and other distortions are ubiquitous in graphical testimony. The perspective that we adopt towards a subject is an expression of judgments about which aspects of that subject are central or important to its characterization and

[11]Why make this claim in terms of Camp's machinery rather than by appealing to more traditional concepts like implicature? One reason is that Camp's tools allow us to be slightly more agnostic with respect to the semantics of graphs. The more important reason is that even if we adopted the implicature terminology instead, we would need something like Camp's concepts in order to capture how non-truth-apt features like the scale affect the speaker's responsibilities. My thanks to an anonymous reviewer and Kenneth Black for pressing me to clarify my thinking here.

which are unimportant or, indeed, entirely irrelevant. By framing a graph in a particular way, therefore, a speaker can communicate that some details or features aren't important or shouldn't be focused on. The contrast between figures 1 and 2a is once again useful here. In both figures, the chosen frame forces a zoomed-out perspective: it highlights the overall trend of the data and downplays the placement of any particular year. In figure 1, however, the years are still individually represented, and while it would be difficult at best to determine whether a particular year was incorrectly placed, the presenter of figure 1 is responsible for the accuracy of each dot. With figure 2a, by contrast, the choice to use a (smoothed) line rather than individual dots eliminates that responsibility: the individual years have disappeared from the picture. What's relevant is the trend line and the trend line alone.

In other words, the framing of a graph plays a role in determining which of its features "count" when evaluating its accuracy. Graphs are not unique in this respect; frames seem to play a similar role in determining what counts as accurate representation more broadly. As Greenberg (2018, 875–76) argues, for instance, the conditions of accurate artistic depiction depend on what he calls the "system of depiction" or what we might intuitively gloss as "style": black and white sketches and color paintings have to be evaluated for accuracy in different ways. The simple account of the last section didn't respect this feature of depictions—its view about what counts as accuracy-relevant content was too flat-footed for the complexities of graphical testimony.

To account for normatively unproblematic distortions, therefore, we need an account on which the speaker who presents a graph is committed only to the accuracy-given-the-operative-perspective of the content, not to its accuracy in some more abstract sense. Since perspectives are flexible and "open-ended," the scope of this commitment will often be messy and open to negotiation—doubly-so because some distortions are non-ideal but forgivable or the best that can be expected in the circumstances. Still, the central point is an essential one: at least for the purposes of testimony, what counts as an accurate graph depends on the perspective and thus on framing devices like the choice of scale.

As I flagged in the last section, this amendment to the simple account can be interpreted in two ways. On the one hand, we might interpret it as an alteration to the semantics of the graph: the "literal" meaning of the graph depends in part on non-truth-apt features like the scale or the choice to use lines instead of dots. On the other hand, we might think that this amendment is best captured by extending the norms of assertion into the pragmatics of graphical testimony. On this reading, the literal meaning of a graph is unaffected by non-truth-apt elements like the scale, but what matters in evaluating the presentation of a graph is its accuracy-given-the-operative-perspective rather than the accuracy of the literal content. Interesting as this question is, however, I want to put it to

the side: regardless of the choice we make here, the result is that the content that the speaker is responsible for depends on the non-truth-apt framing devices that influence the audience's perspective and what matters for present purposes is the speaker's responsibility.

The solution to the second problem follows naturally from the solution to the first. Recall that the second problem was that the appearance of the graph depends in part its on non-truth-apt features and (thus) that scientists often criticize the choices regarding these features using epistemic language. As we've seen, these features are framing devices that play an important role in fixing the perspective, and the perspective is essential to determining how the audience understands the graph and thus what they take away from the testimony. Criticisms of the choice of scale are criticisms of the framing devices. The natural thought, therefore, is that speakers are (in some way) responsible for the framing devices in addition to the truth-apt features of the graph.

Once we notice the important role that framing devices like the scale play in fixing a perspective, it's easy to see why the choice of framing devices would be treated as open to epistemic criticism. Consider an audience presented with figure 1 who comes to believe that climate change is unprecedented because they "see" that it is when they look at the figure. Obviously, this audience relies on the speaker to vouchsafe the accuracy of the data points—the audience's belief is only justified if the data are accurate. But they also exhibit what Fraser (2021) calls "perspectival dependence": if the framing devices force dispositions to "see" what isn't really in the data, then the audience isn't justified in their belief. As such, the audience has to trust the speaker not only with respect to the accuracy of the content but also with respect to the *reliability* of the perspective: the relevant epistemic dispositions should be (sufficiently) truth-conducive in that an audience who views the graph should "see" true conclusions rather than false ones with sufficient consistency. And thus the speaker gains an importantly epistemic responsibility.

There are two important complications here. The first complication is that it's rare for the speaker to be *entirely* responsible for the audience's perspective. In testimonial contexts, perspectives are usually co-determined: even in cases where the speaker has substantial control over the framing of a graph, the audience brings various background assumptions to the testimony and these can affect how they interpret or encounter a graph. In some ways, co-determination limits the speaker's responsibility: she's responsible for her contribution to the audience's perspective—how the content is framed—but not for the audience's contribution. On the other hand, the co-determination of perspectives means that the speaker's contribution should be evaluated in the context of what her audience brings to the perspective. Audiences differ, and the framing choices that will contribute to reliable dispositions in one audience will contribute to unreliable ones in another.

Insofar as the speaker is able to predict the effects of her choices, she's responsible for tailoring her testimony to fit the audience (Dethier, 2022a).

Importantly, tailoring testimony to an audience is not just a matter of fit with their background beliefs—the audience's intellectual goals and interests matter as well.[12] Consider figure 3c. The framing of this graph enforces a perspective that highlights the stability of yearly global averages on a 20°C scale: yearly global averages are not dipping down into permanent winter or shooting up into permanent summer. If the audience's interest lies in the novelty or seriousness of climate change, this perspective is liable to be unreliable, because they're liable to draw false conclusions about the novelty or seriousness of climate change. A different audience with different interests would draw different conclusions, and thus the reliability of a perspective can depend not just on what the audience believes but on what they care about.

Complicating this complication is the fact that audiences are rarely homogeneous.[13] As a consequence, speakers will often have to prioritize fit with the background beliefs and values of some audience members over others. The resulting choices are value-laden (compare Steele, 2013): the choice to prioritize testimonial success with policy-makers (over, say, contrarians or activists) involves value judgments about whose uptake is most important. As such, speakers here face all of the familiar questions from the values-in-science literature, such as who should make these judgments and on what basis. But notice that these (extremely difficult) questions concern the ethical dimensions of testimony, not the epistemic ones. That is, the speaker may be criticized on ethical grounds for prioritizing policy-makers, but this prioritization doesn't betray her specifically epistemic responsibilities. Of course, audience heterogeneity makes (graphical) testimony more difficult, and thus can lead to a familiar gap between cases where testimony is less than fully successful and cases where the speaker is blameworthy. Sometimes, successful testimony is hard and the best that anyone can be expected to do is something that falls short of our evaluative standards.

The second complication is more straightforward: in evaluating whether a perspective is sufficiently truth-conducive, we tend not to weigh all truths equally.[14] Some aspects of what's depicted in a graph are more important than others (of course, this too is a value judgment). In presenting figure 1, for example, we tend to care less about whether the framing allows the audience to draw true conclusions about global temperatures during the middle ages than we do that it allows them to draw true conclusions about the present moment. This is part of what makes smoothing acceptable in examples like figure 2a: while the perspec-

[12]Schroeder (2017, 2022) has previously emphasized this point, though his focus is on ethical constraints on testimony rather than the epistemic commitments of speakers.

[13]My thanks to an anonymous reviewer for pressing me on this point.

[14]My thanks here as well, to both an anonymous reviewer.

tive doesn't allow us to draw true conclusions about individual years with any reliability, it does support truth-conducive inferences with respect to the more important feature, namely the overall trend in the data.

Stepping back: incorporating the changes discussed in this section yields an account on which graphical testimony involves two highly intertwined commitments. First, the speaker commits to the reliability of the framing devices chosen in the context of the salient background assumptions and values. Second, the speaker commits to the accuracy-given-the-operative-perspective of the graph. In our running example, this means that the presenter of figure 1 would be committed first to the scaling (etc.) of the graph being such that you can trust the perspective that is determined by these framing devices and second to the accuracy of the trend and any other features highlighted by the perspective.[15]

What is the nature of this commitment? Here I want to defer: it seems to me that the commitment involved in graphical testimony is just the same as the commitment involved in linguistic testimony. So, if the knowledge norm of assertion captures the epistemic commitment involved in testimony, it's the correct norm for graphical testimony and an agent should present a graph only if she knows that (a) her choices regarding framing devices will contribute to a reliable perspective and (b) the graph's content is accurate from that perspective. If we follow García-Carpintero (2023) and others in defining assertion in terms of the nature of the commitment that the speaker undertakes, then the account that I've offered sees graphical testimony as a kind of assertion—it's an assertion where there's a commitment to both a content and a perspective.

# 3 Beyond graphical testimony

There's an objection that I suspect is likely to have been percolating in the minds of many readers. Plausibly, we should expect testimony to be a (relatively) unified phenomenon, norms-wise: since linguistic and graphical testimony play the same kind of epistemic role, it would be surprising to discover that the two domains are governed by radically different norms. And it may seem like that's the result that

[15] I've presented these as two different commitments, but it's not clear that they're truly separable. As already emphasized, one possible reading of the arguments given so far is that the semantics of graphs depends on non-truth-apt features like the scale. If that is the right conclusion, then the two commitments may well collapse into a single commitment to the literal meaning of the graph being true (modulo the standard context-sensitivity caveats from the assertion literature). And whether these two commitments can come apart has implications for our present discussion: if they can, then our account should say something about what's demanded of an agent when the two commitments aren't mutually satisfiable. Unfortunately, these questions must be left to future work. (My thanks to Claire Dartez for helping me see this point.)

my account delivers. In particular, if my account is right, then—at least given some views common in the testimony literature (e.g. Fricker, 2012)—testifiers are epistemically responsible for much more in the graphical case than they are when using everyday declarative sentences. Indeed, some of the commitments that I've claimed they're responsible for look more like insinuations than traditional assertions.

I think there are a number of reasons not to be worried by this objection. For one thing, the thought that speakers are epistemically responsible only for the "literal" content of their statements and not for insinuations or presuppositions has (rightly) come under attack of late (see Camp, 2018a; Viebahn, 2020).[16] For another, as Camp (2018b, 42) points out in a discussion of maps, the lines between syntax, semantics, and pragmatics are often much blurrier in depictive cases than in the linguistic. Even given the presumption of unity sketched in the last paragraph, therefore, it shouldn't be surprising to find more responsibility for pragmatics in the graphical case given how hard it is to pull the pragmatics of a depiction apart from the semantics.

Neither of these responses is the one that I want to lean on, however. Instead, I want to argue that my account is compatible with these traditional views on assertion—essentially, presenting a graph is like asserting that "this graph is a good representation." If I'm right, not only are the two accounts compatible, we should expect to find similar patterns of commitment in examples beyond the graphical case.

The suggestion that perspectives are relevant to the epistemic responsibilities of a testifier is not entirely new. In a recent paper that has heavily influenced my own thinking on the subject, Rachel Fraser (2021) argues that testimony often fails to accord with what she calls "simple testimony." In particular, she urges that testimony is often *narratively structured*—in the sense of being a story of one kind or another—and that the narrative structures of these stories enforce a perspective that the audience must take on trust. Given Fraser's account of the descriptive practice of narrative testimony, it seems plausible that the norms of narrative testimony would be akin to the norms of graphical testimony: in both cases, there are important framing devices that affect the perspective that the audience adopts.

Similar cases are widespread. Consider explicit simplifications. A concrete example comes from `SkepticalScience.com`, a website dedicated to debunking myths about climate change. At the top of each of their articles, they have a small box labeled "What the science says..." that offers a "basic," "intermediate," or "advanced" summary of the article's content. On the page regarding

[16] See also Lackey (2011), who argues that expert speakers have responsibilities beyond just knowing the literal content of their statements.

climate model accuracy, for instance, the "basic" box says "Models successfully reproduce temperatures since 1900 globally, by land, in the air and the ocean" (Skeptical Science, 2018). This statement is a simplification—as the rest of the article makes clear—and in evaluating it, we should not treat the simplification itself as asserted, because that's (explicitly) not what the authors of Skeptical Science intend.[17]

Instead of the literal truth of the declarative sentence found within the "What the science says..." box, the authors of Skeptical Science are committed to the truth of something like: "What follows is a good summary of the content of this article: Models successfully reproduce temperatures since 1900 globally, by land, in the air and the ocean." Intuitively, a summary is good to the extent that it accurately represents all and only the most important features of whatever is being summarized. Plausibly, then, the epistemic commitments involved in summarizing are similar to those involved in presenting a graph: the speaker is responsible both for identifying the "most important" features and for accurately presenting those features. If the audience is to accept the testimony, they must trust that the speaker has fulfilled their responsibilities in both respects.

This example is representative: scientific communication is shot-through with simplifications and other distortions such as idealizations and approximations. Feminist theorists have been particularly sensitive to this fact: as Grasswick (2010) and Schiebinger (2008) have stressed, scientists engaging in testimony must make decisions about what information to exclude, what to include, and where to simplify or idealize—it's impossible for an expert to communicate everything that they know to an audience with limited background information, time, and attention. In our terminology, these are framing choices: summaries and simplifications rely on tacit assumptions about which features are central and important and which are peripheral. And, of course, framing choices are not epistemically innocent: they affect—often in predictable ways—what the audience comes to believe on the basis of the expert's testimony.

That's not to say that these examples are exactly like the cases of graphical testimony that we examined above. One important difference is that graphs like those examined above make some (but not all) of the framing choices explicit. In figure 1, the audience can see what scale the speaker has chosen. That's largely untrue when a speaker summarizes: while qualifications that mark a simplification such as "roughly" or "basically" are common, it's far less common for speakers to make explicit the details of their simplification—to say what aspects of the target the simplification distorts.

[17] As one reviewer rightly pointed out to me, the quoted claim should probably be read as a generic (for details, see Leslie & Lerner, 2022), in which case the stated claim might be literally true. I don't think granting the point affects my analysis; the main aspect of the example that I care about is that this claim is explicitly flagged as a simplification.

As a consequence, the perspectives employed in linguistic testimony are often *presupposed* rather than *presented* as part of the testimony itself. It seems to me that this doesn't negate the speaker's responsibility for the choice of framing: the audience has to trust her with respect to those choices regardless (compare Viebahn, 2020). Figure 1 relies on particular framing devices, and the audience who "sees" that climate change is unprecedented has to trust that framing if they are to justify accepting what they see. But the same is true of the bare statement "climate change is unprecedented." The framing choices—regarding, for instance, the timescales and features involved when evaluating "unprecedented" (Watkins, 2023)—are simply presupposed rather than built into the graph.

This difference in the explicitness of the perspective isn't a difference in kind, therefore, because important choices about the nature of the perspective can be made explicit in the linguistic case or presupposed in the graphical one. Nevertheless, it's worth emphasizing, because it's crucial to recognize just how important the choice of perspective is in expert testimony and how minimally these choices may be indicated or flagged. As both Camp (2017, 2018a) and Fraser (2021) emphasize, presuppositions can "force" a particular perspective on an audience, and this effect is particularly easy to achieve in the context of expert testimony, where audiences are liable to be ignorant of the the nature of framing devices employed, the effects of a particular framing choice, the alternative framings that could be adopted, and the reasons that favor one framing over another. The audience who accepts a summary like the one offered by Skeptical Science is likely to have very little understanding of the perspective under which the summary is a good one, but they are nevertheless relying on that perspective when they subsequently rely on the summary itself.

If the foregoing is broadly right, then at least some of expert testimony more broadly is perspectival in much the same way as graphical testimony is: experts consistently take on responsibility for both framing and content, and their responsibility for the latter must be evaluated in the context of their commitment to the former. This responsibility comes by way of taking responsibility for a simplification or summary—the view I've been outlining, in essence, is that being responsible for a claim like "to summarize, ..." is just to be responsible for both choices regarding framing and the perspective-relative content.

This phenomenon provides both a rejoinder to the objection outlined at the beginning of the section and a compelling explanation for why we find the pattern of responsibilities that we do in the graphical case. Regarding the rejoinder, not only is the (apparent) gap between linguistic and graphical testimony not *that* surprising, the same pattern of responsibilities can be found in other cases including more familiar cases of linguistic testimony. Regarding the explanation, we can think of presenting a graph in the appropriate testimonial setting as akin to asserting that the graph is a "good representation," where "good represen-

tation" is to be read in much the same way as "good summary" was above. If we accept this explanation, it would further explain why the norm that governs linguistic testimony (whether knowledge, justification, belief, etc.) applies to the graphical case in the way that I've claimed it does, namely that there's a sense in which the responsibilities of graphical testimony can be thought of as inherited from the tacit assertion that the graph is a good representation.

Allow me to step back. It is unsurprising that linguistic and graphical testimony should be unified in *some* way. This principle alone suggests that if (e.g.) the knowledge norm governs assertion, then a presenter must know something about a graph to properly "assert" it. The question raised in this paper is essentially: *what* must she know about the graph? As we've just seen, there is a class of linguistic testimony about a representation (such as a graph or picture) in which the speaker commits to claims about both truth-apt and non-truth-apt features of said representation. My contention in this section is that presenting a graph is akin to linguistic testimony of this sort: when a speaker testifies by way of a graph, she is always "saying" something both about the truth-apt and non-truth-apt features of the graph. *Proper* graphical presentation thus requires knowledge (or belief, justified belief, etc.) about both sorts of features. Specifically: the speaker must know that (a) the graph's framing devices are reliable and (b) the graph's content is accurate relative to the operative perspective.

Here's another way of approaching the point. Unlike linguistic testimony, graphical testimony always involves the presentation of non-truth-apt features like the scale. A major question that any account of graphical testimony should answer, therefore, is how these features affect the propriety of the testimony: how do we get from the choice of scale to proper or improper presentation? Notably, this question is not answered by assuming that a particular norm of assertion applies in the graphical case. What bridges the gap and allows us to answer this important question are the concepts of frame and perspective that we've borrowed from Camp's work. And what we've seen in this section is that we can use these tools to answer this question without undermining the assumption of the unity. Indeed, we've seen that there are reasons to expect that these tools will be helpful in a range of linguistic cases as well—specifically, those cases in which the content of the assertion involves both truth-apt and non-truth-apt features of a representation.

# 4 Perspectives and honest testimony

Much of the literature on scientific testimony is focused on "values." One particular locus is the question of whether and when scientists and other experts should exhibit communicative virtues such as honesty. Some authors take honesty to be

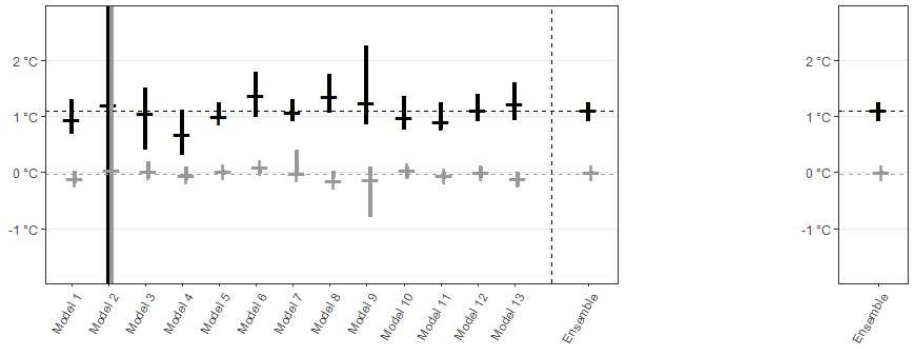


(a) Full graph (b) Truncated graph

Figure 4: Two graphs representing the human contribution to warming. Black lines are estimates of the human contribution; grey lines, the natural contribution. Graph (a) is patterned on similar graphs found throughout climate science (see, e.g., IPCC, 2021, 439, fig. 3.7c). Graph (b) is "truncated" in the sense that it includes only the composite "ensemble" estimate from graph (a). Illustrative only; while realistic, not constructed using real data.

a side constraint on testimony: there are many permissible choices that one can make as an expert witness, but dishonest testimony is impermissible (Keohane, Lane, & Oppenheimer, 2014; Nisbet, 2009; Schroeder, 2017). Insofar as these authors are right, one of the most important questions for an account of graphical testimony is "when is it honest?" The aim of this section is thus to show how my account answers this question.

At the same time, some philosophers (e.g. John, 2018) have argued that there is at least an apparent conflict between the widespread view that dishonesty is impermissible and the truism that both science and scientific testimony are shot-through with distortions: certainly, there's a conflict between idealized or simplified testimony and any understanding of honesty that equates it to telling "the truth, the whole truth, and nothing but the truth" (see, e.g., Irzik & Kurtulmus, 2019). An appealing consequence of my view is that it shows why this apparent conflict is in fact illusory, at least in the graphical case: properly understood, honesty is compatible with even tacit distortions.

To begin, consider figure 4. Graphs like figure 4a are often used to represent different estimates of the human contribution to warming. Of course, figure 4a doesn't capture the truth, the whole truth, and nothing but the truth with respect to the causes of global warming. I'm inclined to argue that *no* graphical depiction of climate change does that; there are simply too many relevant variables for "the whole truth" to be presented in two dimensions. Still, there's an important sense

in which figure 4b is "further" from the whole truth than figure 4a is: it includes less potentially relevant information about the current state of the evidence. And so the contrast between the two graphs provides a useful example for our goal of understanding what "honesty" means in the context of graphical testimony.

Recall that on the account defended in this paper, the speaker who presents one of the graphs in figure 4 is committed to two things: the in-context reliability of the framing devices employed and the perspective-relative accuracy of the content. There's (substantial) debate in the literature about the necessary conditions on honesty in standard linguistic cases (Mahon, 2015), but it's widely agreed that a speaker is dishonest if they knowingly present a falsehood. The natural extension to the graphical case would involve the speaker presenting a depiction where they know either (a) that the relevant framing choices are unreliable or (b) that the content is inaccurate (when viewed from the relevant perspective).

The second of these conditions is familiar and straightforward: honest testimony requires, at minimum, that the evaluation-relevant content is not known to be false. The only element of this condition that is new in any way is that we've restricted what counts as "evaluation-relative" to account for the perspectival nature of graphical testimony. The first condition is less familiar. On its face, it appears simple: to be honest, graphical testimony cannot be framed in such a manner that the resulting perspectives are (known to be) unreliable. I cannot honestly present figure 4b in a setting where the question under discussion (Roberts, 2012) concerns the range of results offered by different models without changing the question.

As I argued in section 2, the reliability of a perspective is a function of the background beliefs and values of the audience. As a consequence, whether an act of graphical testimony is honest will depend on the audience as well. That's not to say that honest testimony can be rendered dishonest merely by altering the setting—our ethical judgments about testimony are more complex than that—but it is to say that we cannot ignore the assumptions and values of the audience when evaluating a given piece of testimony.

To illustrate this point, it will be helpful to detour into the arguments that Stephen John (2018) offers for his claim that there is a conflict between honesty and the practices of idealization and simplification that are common in the sciences. Drawing from Wendy Parker—particularly Parker (2014), though see also Parker (2010) and Parker & Risbey (2015)—John offers the following case:

> Assume, then, that a climate scientist knows that she could report a single precise probability estimate to policy-makers, and doing so would secure that policy-maker's action against climate change. However, she cannot (in Parker's terms) 'own' that prediction; at best, she can 'offer' it because she is aware that her estimate is subject to

> significant 'second-order uncertainty'. Unfortunately, reporting these uncertainties would be more likely to lead to inaction. I presume a proponent of 'honesty' thinks that the scientist should, still, communicate the less precise estimate because doing otherwise is 'dishonest', in the sense that the scientist would be representing the precise estimate as enjoying a privileged status which it lacks. (John, 2018, 83)

One of the implications of my account is that—at least in the setting of graphical testimony—John's "presumption" is wrong: what the proponent of honesty thinks the scientist should do depends on details that haven't been established.

Following John, let's assume that the imprecise estimate is in some important sense a more "accurate" representation of the true uncertainty than the precise estimate.[18] Plausibly, this implies that figure 4b is less accurate than figure 4a. Technically, the error bars in these graphs are constructed using a single precise probability distribution: since the methods employed in this area are largely classical, we can assign a p-value of precisely .9 to the range spanned by the error bar. So if *precision* is the problem, then figure 4b is objectionable on the letter of the critique. And the difference between figures 4a and 4b is certainly in the spirit of the critique: relative to figure 4a, figure 4b hides the "second-order uncertainty" present in these estimates by hiding (e.g.) that "Model 2" fails to deliver well-behaved estimates at all.

Is granting the assumption that the first figure offers a more accurate representation sufficient for testimony involving the the second to be dishonest? No. Consider the following case that follows John's formula:

> Cooperation. Marie is a climate scientist; she works with Doug, a policy-maker, in developing responses to climate change. In their interactions, Marie and Doug aim solely to develop the best possible policies (whatever those are). Marie knows that imprecise probabilities provide a more accurate representation of the evidence than precise ones. She also knows (a) that Doug lacks training in statistics, economics, or formal epistemology—he's not familiar or comfortable with imprecise probabilities—and (b) that the policy choices available aren't sensitive to small differences in probability. Reasoning that presenting the imprecise probability would only serve to confuse Doug and waste both of their time, Marie opts to present Doug with figure 4b, which hides the 'second-order uncertainty' and presents estimates based on a single precise probability.

[18]The idea that there's something inaccurate about precise probabilities in the setting of climate science is contentious. See Dethier (2022b) and Katzav et al. (2021) for contrasting opinions.

Intuitively, it seems to me that Marie is being honest; my account delivers the same result.

According to the analysis that I've offered over the last couple sections, Marie is committed to two things: the in-context reliability of the framing devices and the perspective-relative accuracy of the graph's content. Whether she has fulfilled these commitments depends on the details of the case. We stipulated that the imprecise estimate is in some important sense a more "accurate" representation of the true uncertainty than the precise estimate, but that doesn't determine the issue. A color photograph is in some sense a more accurate representation of the colors of a subject than a black and white sketch is, and yet presenting the latter is not necessarily dishonest. The question that we have to ask in this situation is whether Marie is right to deem the difference between the precise and imprecise probabilities unimportant given Doug's epistemic interests. Is Doug liable to draw false conclusions due to the framing that Marie has adopted? Or are Doug's assumptions and goals such that this way of simplifying is unproblematic? The latter situation seems plausible: while (by stipulation) the imprecise probabilities are more accurate in some abstract sense, they're less well-suited to providing Doug with the knowledge he wants and needs.

Suppose, therefore, that the framing devices Marie employs in simplifying are reliable, that the precise probabilities are accurate from within the resulting perspective, and that Marie knows both these facts. There's then nothing about honesty as traditionally understood that would render Marie's presentation inappropriate. After all, Marie has no intention of deceiving or misleading Doug, nor does she have any reason to believe that she will. And it's not as though Doug asked her for the truth, the whole truth, and nothing but the truth. Instead, we stipulated that Doug's sole aim is to develop the best possible response to climate change—while information beyond what's necessary for his policy-making purposes might be part of "the whole truth," it isn't part of what he asked for. (It isn't the "question under discussion.") Since Marie's judgment is that presenting the precise probabilities is the best way of achieving their shared aim, we might criticize her for incorrectly evaluating the situation, but we can't criticize her for being dishonest, at least in any traditional sense: she intends neither to lie nor to deceive. This case thus provides us with an example of an honest distortion, or at least a sketch of what honest distortion might look like.

Contrast Cooperation with the following:

> Conflict. Amir is a climate scientist; he works with Linda, a policy-maker, in developing responses to climate change. In their interactions, Amir aims solely to develop the best possible policies. By contrast, Linda receives large campaign donations from the oil lobby, and her sole aim is to develop policies that will increase revenues for her

> donors. Amir knows that imprecise probabilities provide a more accurate representation of the evidence than precise ones. He also knows that Linda will jump on any opportunity to delay action that would decrease company profits, and that the imprecision of imprecise probabilities would serve exactly that purpose. Reasoning that a precise probability leaves Linda with fewer options for delaying action, Amir opts to present Linda with figure 4b, which hides the 'second-order uncertainty' and presents estimates based on a single precise probability.

It seems to me that Amir is being dishonest.

The analysis supports this intuition. Unlike Doug, Linda is interested in facts that go beyond those that facilitate making the best policy—indeed, she is only interested in facts that she can use to support her preference for inaction. As such, a perspective that downplays or ignores higher-order uncertainty doesn't align with her epistemic interests in the same way it does with Doug's and, thus, she's liable to form an inaccurate picture of the state of the science on the basis of this presentation. In particular, she's liable to conclude that there's nothing in the science that she can use to achieve her goal of increasing revenue for her donors. Importantly, Amir knows this about Linda and has nevertheless chosen to present a perspective that downplays the information that she thinks is important and is likely to focus on. He may not have lied to Linda, but he's been dishonest by (intentionally) misleading her.[19]

These examples are obviously stylized, but the contrast between them nicely illustrates both how the interests of the audience can affect the honesty of a presentation and the possibility of honest distortions. Traditionally, honesty has been understood as a matter of refraining from intentionally misleading the audience. In the context of graphical testimony, avoiding misleading the audience is a matter of choosing both the right content and the right perspective from which to view that content. In this framework, simplifications and other distortions are perfectly honest so long as the chosen perspective is one on which the distorted elements are unimportant. Not everything goes, of course: as figure 2b illustrates, even subtle distortions can be systematic and misleading enough to be impermissible. But distortion (or "inaccuracy") in an abstract sense is not inherently a barrier to honesty: what matters is how the inaccuracy figures into the purposes of the speaker and audience.

[19] A potential complication here is that at least some people have the intuition that it would be more dishonest for Amir to assert that the graph is a good one. I don't share that intuition, but potential differences in our reactions to linguistic and non-linguistic testimony deserve more attention. See Lewerentz & Viebahn (2023) for an account of pictoral assertion that might be extrapolated to give the opposite reading in this case.

Our earlier concern about heterogeneous audiences returns at this juncture. It is one thing to say that a presentation (or assertion, in the linguistic case) is less than perfectly successful because the speaker chose to prioritize successful communication with some audience members over others. It's another thing entirely to say that it is *dishonest*. Here too, however, I'm inclined to say that we can disentangle the potential cases. Even the most virtuous speakers must run the risk of misleading some individuals when presenting to (unpredictably) heterogeneous audiences. I am disinclined to call these cases of dishonesty. But misleading a subset of the audience can also be an intended feature of a graph. Figure 3b, for example, could be presented with the intent of deceiving uncareful readers, particularly if accompanied by what Jennifer Saul (2023) calls a "figleaf," namely the excuse that the data are in fact correctly plotted.[20] Such cases are clearly dishonest.

Between these two extremes are cases of *predictably* but not intentionally misleading a subset of the audience. Here I suspect that our intuitions will depend heavily on feasibility judgments: there's something dishonest about choosing a representation that misleads some of your audience when you could easily have chosen one that would not have had that effect. My account is largely silent on how we should thread our way through the moderate cases, however. I think this is as it should be. What my account makes clear is that whether graphical testimony is honest depends on choices regarding the framing and the perspective-relative content. But most of the work in dividing honest cases from dishonest ones ought to be done by the details of an account of honesty, not by our account of testimony.[21]

One other important lesson to take from our analysis is that dishonesty is a product of conversational situations where the speaker and audience find themselves (or perceive themselves as) at cross-purposes (compare Camp, 2018a). When the speaker and audience are perfectly aligned—when, as in Cooperation, the speaker and audience have the same epistemic aims for the conversation—there's no motivation for dishonesty. Or, more simply, when the speaker and audience desire the same outcomes, it's easy to find a shared perspective that allows for distortions that don't mislead (or don't seriously mislead) the audience. When the speaker and audience aren't perfectly aligned in their epistemic aims for the

[20]Similarly, depictions such as choropleth maps of presidential vote winners by county can be used as the kind of dishonest speech known as "dogwhistles"; again, see, Saul (2023).

[21]That said, notice that while a speaker can successfully testify with respect to some audience members and fail to do so with respect to others, there doesn't seem to be a sense in which they can be honest with respect to some and dishonest with respect to others. That is, honesty seems to necessarily apply to the speech-act as a whole, whereas success can be evaluated on an audience-member by audience-member basis. So audience heterogeneity does not imply the existence of testimony that is simultaneously both honest and not.

conversation, by contrast, there will be no perspective that respects the values of all parties: they disagree about what features of the target are important (central, relevant, etc.). In these cases, distortions that mislead no members of the audience will be much harder to find.

Importantly, the divergence between speaker and audience need not be large to motivate dishonesty (O'Neill, 2002). For instance: many people are under the false impression that looking at the "raw data" is the best way to determine the right policy decisions. Even where the speaker and audience are aligned with respect to their ultimate goals in a conversation, therefore, the speaker may be motivated to dishonesty by the audience's mistaken beliefs about what constitutes good evidence. This consequence might seem like a problem for my view, an indication that the view is overly demanding in one respect or another. I think this is the wrong diagnosis, however. My view is that honesty is *at least* an ideal of good testimonial practice, and that honesty is demanding. But good testimonial practice is not always obligatory: there are cases where dishonest testimony is justified in the name of the greater good. Perhaps Amir is doing the right thing. But if so, the best way to capture that fact is in terms of his case being one where the ethical reasons that favor being a good (read: honest) communicator are outweighed by those that favor misleading his audience.

Finally, there's a neat explanation for why distortions and honesty are compatible in the context of graphical testimony. As argued in the last section, presenting a graph is akin to asserting that the graph is a good representation, and—of course—a representation can be good while being simplified, abstracted, idealized, etc. So insofar as the distorted graph is actually good—insofar as the perspective is reliable and the content accurate from that perspective—the presenter has managed to present something that is simplified in a perfectly honest way. Notice, importantly, that this neat explanation suggests that distortions are compatible with honesty in some—but not all—linguistic cases as well. What's crucial is that the distorted claims are not asserted directly and are instead presented as a kind of representation (e.g., a summary) that allows for distortion.

The takeaway is as follows: a speaker can honestly present a simplified or idealized graph so long as that graph is appropriately framed. The appropriate framing is (relatively) easy to find when there's an alignment of interests and values between speaker and audience, and hard—if not impossible—to find in cases where the interests and values of the speaker and audience diverge and/or where the audience is sufficiently heterogeneous. In the latter cases, it will be much more difficult to avoid misleading at least some audience members, but even in these cases, something more seems to be required for dishonesty. After all, it seems in-apt to judge someone "dishonst" when they have done everything in their power to communicate the truth. And once we recognize that mis-aligned interests and audience heterogeneity make successful testimony difficult, it should

quickly become clear that in many cases, experts who are doing their best may still fall short of full success.

# 5 Conclusion

Much of scientific testimony is communicated via depictions such as graphs. Graphs have both truth-apt and non-truth-apt features, but the latter influence the audience's perspective towards the former in a wide variety of ways. This influence affects the epistemic responsibilities of the speaker: not only is she now epistemically responsible for how her choices influence the audience's perspective, her responsibility for the content must be evaluated with the relevant perspective in mind. I've argued that the interplay between these responsibilities makes room for honest distortions, and that much more of expert testimony is perspectival in this same respect: neither experts nor audiences have the time for "the whole truth" about a subject and testimony thus often involves adopting a partial perspective towards the content.


# Acknowledgments

My thanks to Liam Kofi Bright, Claire Dartez, Mikkel Gerken, Stephen John, three anonymous reviewers at Noûs, and audiences in Bielefeld, Bristol, Edinburgh, Minneapolis, and Portland for comments on earlier versions of this paper. My thanks also to Elisabeth Camp for an extremely helpful conversation on her work and its application to graphs and other depictions.

# Funding

Funding for this research was provided in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) Project 254954344/GRK2073 and in part by the National Science Foundation under Grant No. 2042366.


# References


Azzouni, J. (2013). The Relationship of Derivations in Artificial Languages to Ordinary Rigorous Mathematical Proof. *Philosophia Mathematica* 21: 247–54. DOI: 10.1093/philmat/nkt007.

Bach, K. (2008). Applying Pragmatics to Epistemology. *Philosophical Issues* 18: 68–88. DOI: 10.1111/j.1533-6077.2008.00138.x.

Buckwalter, W. (2023). The Belief Norm in Academic Publishing. *Ergo* 9: 1485–1507. DOI: 10.3998/ergo.3117.

Camp, E. (2017). Why Metaphors make Good Insults: Perspectives, Presupposition, and Pragmatics. *Philosophical Studies* 174: 47–64. DOI: 10.1007/s11098-015-0525-y.

— (2018a). Insinuation, Common Ground, and the Conversational Record. In: *New Work on Speech Acts*. Ed. by D. Fogal, D. W. Harris, & M. Moss. Oxford: Oxford University Press: 40–66.

— (2018b). Why Maps Are Not Propositional. In: *Non-Propositional Intentionality*. Ed. by A. Grzankowski & M. Montague. Oxford: Oxford University Press: 19–45.

— (2019). Perspectives and Frames in Pursuit of Ultimate Understanding. In: *Varieties of Understanding: New Perspectives from Philosophy, Psychology, and Theology*. Ed. by S. Grimm. Oxford: Oxford University Press: 17–45.

Dang, H. & L. K. Bright (2021). Scientific Conclusions Need Not Be Accurate, Justified, or Believed by their Authors. *Synthese* 199: 8187–203. DOI: 10.1007/s11229-021-03158-9.

de Toffoli, S. (2017). 'Chasing' the Diagram—The use of Visualizations in Algebraic Reasoning. *The Review of Symbolic Logic* 10: 158–86. DOI: 10.1017/s1755020316000277.

Dethier, C. (2022a). Science, Assertion, and the Common Ground. *Synthese* 200: 1–19. DOI: 10.1007/s11229-022-03580-7.

— (2022b). When is an Ensemble Like a Sample? 'Model-Based' Inferences in Climate Modeling. *Synthese* 200: 1–20. DOI: 10.1007/s11229-022-03477-5.

Fleisher, W. (2021). Endorsement and Assertion. *Noûs* 55: 363–84. DOI: 10.1111/nous.12315.

Fraser, R. (2021). Narrative Testimony. *Philosophical Studies* 178: 4025–52. DOI: 10.1007/s11098-021-01635-y.

Fricker, E. (2012). Stating and Insinuating. *Aristotelian Society Supplementary Volume* 86: 61–94. DOI: 10.1111/j.1467-8349.2012.00208.x.

García-Carpintero, M. (2023). Lying versus Misleading, with Language and Pictures: The Adverbial Account. *Linguistics and Philosophy* 46: 509–32. DOI: 10.1007/s10988-022-09355-0.

Gerken, M. (2022). *Scientific Testimony: Its roles in Science and Society*. Oxford: Oxford University Press.

Giere, R. N. (2006). *Scientific Perspectivism*. Chicago: University of Chicago Press.

Grasswick, H. E. (2010). Scientific and Lay Communities: Earning Epistemic Trust through Knowledge Sharing. *Synthese* 177: 387–409. DOI: 10.1007/s11229-010-9789-0.

Greenberg, G. (2018). Content and Target in Pictorial Representation. *Ergo* 5: 865–98. DOI: 10.3998/ergo.12405314.0005.033.

— (2021). Semantics of Pictorial Space. *Review of Philosophy and Psychology* 1: 847–87. DOI: 10.1007/s13164-020-00513-6.

IPCC (2021). *Climate Change 2021: The Physical Science Basis*. Ed. by V. Masson-Delmotte et al. Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge: Cambridge University Press.

Irzik, G. & F. Kurtulmus (2019). What Is Epistemic Public Trust in Science? *British Journal for the Philosophy of Science* 70: 1145–66. DOI: 10.1093/bjps/axy007.

John, S. (2018). Epistemic Trust and the Ethics of Science Communication: Against Transparency, Openness, Sincerity and Honesty. *Social Epistemology* 32: 75–87. DOI: 10.1080/02691728.2017.1410864.

Katzav, J. et al. (2021). On the Appropriate and Inappropriate Uses of Probability Distributions in Climate Projections, and Some Alternatives. *Climatic Change* 169: 1–20. DOI: 10.1007/s10584-021-03267-x.

Keohane, R. O., M. Lane, & M. Oppenheimer (2014). The Ethics of Scientific Communication under Uncertainty. *Politics, Philosophy and Economics* 13: 343–68. DOI: 10.1177/1470594X14538570.

Kosslyn, S. M. (2006). *Graph Design for the Eye and Mind*. Oxford: Oxford University Press.

Lackey, J. (2008). *Learning from Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.

— (2011). Assertion and Isolated Second-hand Knowledge. In: *Assertion: New Philosophical Essays*. Ed. by J. Brown & H. Cappelen. Oxford: Oxford University Press: 251–75.

Leslie, S.-J. & A. Lerner (2022). Generic Generalizations. In: *Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. URL: https://plato.stanford.edu/entries/generics/ (visited on 04/28/2024).

Lewerentz, L. & E. Viebahn (2023). Truth and Directness in Pictorial Assertion. *Linguistics and Philosophy* 46: 1441–65. DOI: 10.1007/s10988-023-09392-3.

Mahon, J. E. (2015). The Definition of Lying and Deception. In: *Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. URL: https://plato.stanford.edu/entries/lying-definition (visited on 09/14/2023).

Mann, M. E. (2013). *The Hockey Stick and the Climate Wars: Dispatches from the Front Lines*. New York: Columbia University Press.

Mann, M. E., R. S. Bradley, & M. K. Hughes (1999). Northern Hemisphere Temperatures During the Past Millennium: Inferences, Uncertainties, and Limitations. *Geophysical Research Letters* 6: 759–62. DOI: 10.1029/1999GL900070.

Massimi, M. (2022). *Perspectival Realism*. Oxford: Oxford University Press.

Nisbet, M. C. (2009). The Ethics of Framing Science. In: *Communicating Biological Sciences: Ethical and Metaphorical Dimensions*. Ed. by B. Nerlich, R. Elliott, & B. Larson. New York: Routledge: 51–74.
O'Neill, O. (2002). *A Question of Trust*. Cambridge: Cambridge University Press.
PAGES 2k Consortium (2019). Consistent Multidecadal Variability in Global Temperature Reconstructions and Simulations over the Common Era. *Nature Geoscience* 12: 643–49. DOI: 10.1038/s41561-019-0400-0.
Pagin, P. & N. Marsili (2021). Assertion. In: *Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta. URL: https://plato.stanford.edu/entries/assertion.
Parker, W. S. (2010). Whose Probabilities? Predicting Climate Change with Ensembles of Models. *Philosophy of Science* 77: 985–97. DOI: 10.1086/656815.
— (2014). Values and Uncertainties in Climate Prediction, Revisited. *Studies in History and Philosophy of Science Part A* 46: 24–30. DOI: 10.1016/j.shpsa.2013.11.003.
Parker, W. S. & J. S. Risbey (2015). False Precision, Surprise and Improved Uncertainty Assessment. *Philosophical Transactions of the Royal Society Part A* 373: 20140453. DOI: 10.1098/rsta.2014.0453.
Perini, L. (2005). The Truth in Pictures. *Philosophy of Science* 72: 262–85. DOI: 10.1086/426852.
Roberts, C. (2012). Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics. *Semantics & Pragmatics* 5: 1–69.
Saul, J. M. (2023). *Dogwhistles and Figleaves: How Manipulative Language Spreads Racism and Falsehood*. Oxford: Oxford University Press.
Schiebinger, L. (2008). West Indian Abortifacients and the Making of Ignorance. In: *Agnotology: The Making and Unmaking of Ignorance*. Ed. by R. N. Proctor & L. Schiebinger. Stanford: Stanford University Press: 149–62.
Schmidt, G. A. (2016). *Comparing Models to the Satellite Datasets*. URL: https://www.realclimate.org/index.php/archives/2016/05/comparing-models-to-the-satellite-datasets/ (visited on 04/02/2023).
Schroeder, S. A. (2017). Using Democratic Values in Science: An Objection and Response. *Philosophy of Science* 84: 1044–54. DOI: 10.1086/694006.
— (2022). An Ethical Framework for Presenting Scientific Results to Policy-Makers. *Kennedy Institute of Ethics Journal* 32: 33–67. DOI: 10.1353/ken.2022.0002.
Skeptical Science (2018). *How reliable are climate models?* URL: https://skepticalscience.com/climate-models.htm (visited on 03/25/2023).
Steele, K. (2013). The Scientist qua Policy Advisor makes Value Judgments. *Philosophy of Science* 79: 893–904. DOI: 10.1086/667842.
Tufte, E. R. (2001). *The Visual Display of Quantitative Information*. 2nd ed. Cheshire: Graphics Press.

van Fraassen, B. C. (2008). *Scientific Representation: Paradoxes of Perspective*. Oxford: Oxford University Press.

Viebahn, E. (2020). Lying with Presuppositions. *Noûs* 54: 731–51. DOI: 10.1111/nous.12282.

Watkins, A. (2023). Scaling Procedures in Climate Science: Using Temporal Scaling to Identify a Paleoclimate Analogue? *Studies in History and Philosophy of Science Part A* 102: 31–44. DOI: 10.1016/j.shpsa.2023.09.004.

Weiner, M. (2005). Must we Know What we Say? *Philosophical Issues* 14: 227–51. DOI: 10.1215/00318108-114-2-227.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.