Goodness

Bart Geurts

If it wasn't obvious that equivalent descriptions may cause differential evaluations, there is a considerable body of experimental evidence to prove the point. For instance, Levin (1987) asked participants to evaluate the hypothetical purchase of ground beef that was described as "75% lean" for one group and "25% fat" for another. Despite the fact that these descriptions are truth-conditionally equivalent (75% lean ground beef is 25% fat, and vice versa), Levin found that the first group produced higher ratings on several scales, including high/low quality and good/bad taste; these effects persisted, albeit at attenuated levels, even after the ground beef had been tasted (Levin and Gaeth 1988). Similarly, when medical treatments were alternatively described in terms of survival and mortality rates (McNeill et al. 1982, Levin et al. 1988) or when research and development teams were alternatively presented in terms of their success and failure rates (Duchon et al. 1989), positive descriptions prompted higher rates of positive responses. In this paper, I study a number of puzzles which were inspired by these experimental findings. I will argue that these puzzles, which are about the interpretation of evaluative statements, call for a novel kind of pragmatic treatment, which I will develop in some detail. Possible connections between this analysis and the experimental data are discussed in Geurts (2010).

Sad tidings. An airplane carrying 600 passengers has crashed in the Pyrenees. 400 people died in the accident; 200 survived. Hence, in this context, the propositions "200 people survived" and "400 people died" would seem to be equivalent. Now consider the following pair of statements:

(1) a. It's good that 200 people survived.
    b. It's good that 400 people died.

According to my intuition, we would tend to read these statements as contradicting each other. It would be decidedly odd for someone who has just uttered (1a) to go on uttering (1b). And this is not just because (1b) is a peculiar statement in its own right, for someone who, depravedly, stated (1b) with full conviction would not be expected to endorse (1a) as well. It is also relevant to note that (1a) and (1b) don't *have* to be construed as contradictories. For the true Panglossian, everything is equally right and good, and therefore one of that tribe could endorse both statements without fear of contradicting himself. There may be a strong preference for a contradictory construal, but it is not mandatory.

It will be obvious, I trust, what my first question is going to be: How is it possible for (1a) and (1b) to be interpreted as contradictories, given that, by hypothesis, their embedded clauses are truth-conditionally equivalent? The second puzzle is the obverse of the first one: How is it possible for (1a) and (1b) to be consistent with (2a) and (2b), respectively?

(2) a. It's bad that 400 people died.
    b. It's bad that 200 people survived.

Again, it would be rather nasty for someone to say (2b), and in this sense the sentence is odd, but that is beside the point. The problem I'm interested in is how (2b) manages to be consistent with (1b). Ditto for (1a) and (2a).

It might be thought that both of these problems admit of a straightforward solution. For it is obvious that, out of context, the sentences "400 people died" and "200 people survived" express distinct propositions. However, these propositions come apart only in worlds in which the number of passengers does not equal 600, and it is hard to see why that should be relevant. Hence, I will stick to the assumption that, in our examples, these sentences are truth-conditionally equivalent.

Now change the scenario somewhat. The number of passengers remains the same, but the exact number of casualties is not yet known. All we have to go on is that more than 200 people survived the crash, or equivalently, that fewer than 400 died. Now consider:

(3) a. It's good that more than 200 people survived.
    b. It's good that fewer than 400 people died.

Unlike (1a,b), this pair is clearly consistent. In fact, (3a) and (3b) would seem to be synonymous. How is this possible? That is my third and last puzzle.

Although I won't be able, on this occasion, to completely solve all three puzzles, I do believe I can offer the outlines of a plausible solution. To explain the guiding idea, let me show how (1a) and (1b) might come to contradict each other. On the one hand, if someone utters (1a), we tend to infer that, according to the speaker:

(4) It would have been better if more than 200 people had survived and worse if fewer than 200 people had survived.

On the other hand, if someone uttered (1b) with sufficient conviction, we would be inclined to infer that, according to the speaker:

(5) It would have been better if more than 400 people had died and worse if fewer than 400 people had died.

These inferences are incompatible, and that's why (1a) and (1b) are contradictories on what I take to be their most natural readings. Hence, the key idea is that the speaker's evaluation of an actual state of affairs may carry information about how he would have evaluated alternative states of affairs. The main goal of this paper is to investigate the mechanism underlying such counterfactual implications, and to show in more detail how they will help to solve our three puzzles.

The inference in (4), for example, would be accounted for if we could assume that whatever quality is expressed by "good" is positively correlated with the quantitative scale on which the embedded clause of (1a) is sitting. Let me explain. Since "good" is a gradable adjective, its interpretation is relative to a comparison set. For example, if I say "It's good that it's raining", the comparison set might simply be $\{[\![\text{it's raining}]\!], [\![\text{it's not raining}]\!]\}$, in which case my utterance implies that the first is better than the second, or more formally: $g([\![\text{it's raining}]\!]) > g([\![\text{it's not raining}]\!])$, where g is a "goodness function", which maps propositions onto qualitative degrees.

When "good" combines with a quantifying proposition like $[\![\text{200 people survived}]\!]$, the comparison set might be $\{[\![n \text{ people survived}]\!] \mid 0 \leq n \leq 600\}$. Besides being ordered in qualitative terms, this set also comes with a quantitative ordering, which I will symbolise by "$\preceq$":

(6) $[\![\text{0 people survived}]\!] \prec [\![\text{1 person survived}]\!] \prec [\![\text{2 people survived}]\!] \ldots$

I assume that $\preceq$ may but not need not be an entailment ordering; in the current example it isn't.

Now, we can capture the inferences in (4) and (5) as follows:

(7) *Co-optation* (strong version)
$\forall \varphi, \psi$: if $\varphi \succ \psi$, then $g(\varphi) > g(\psi)$.

The label "co-optation" derives from the intuition that the quantitative ordering $\preceq$ is in a sense co-opted for fleshing out the qualitative ordering induced by g. I will have more to say about this presently. As defined in

(7), co-optation is a rather strong assumption to make, but it should be noted that the most obvious way of weakening it will render it too weak:

(8) *Co-optation* (weak version)
$\forall\varphi, \psi$: if $\varphi \succeq \psi$, then $g(\varphi) \geq g(\psi)$.

While the strong version of co-optation says that more of a good thing is better (and more of bad thing is worse), the weak version merely entails that more of a good thing is not worse (and more of a bad thing is not better), which is too weak for deriving the inferences in (4) and (5).

The following version of co-optation is strictly weaker than (7), but slightly stronger than (8), and strong enough for our purposes:

(9) *Co-optation* (medium-strong version)
a. $\forall\varphi, \psi$: if $\varphi \succeq \psi$, then $g(\varphi) \geq g(\psi)$, and
b. $\exists\varphi, \psi$: $\varphi \succ \psi$ and $g(\varphi) > g(\psi)$.

This says that, as you go down a series of propositions lined up by increasing strength, goodness never decreases and increases at least once. Let's apply this to our first puzzle:

(10) a. It's good that 200 people survived.
b. It's good that 400 people died.

According to (9a), it follows from (10a) that

(11) $\forall m \geq n$: $g([\![m \text{ people survived}]\!]) \geq g([\![n \text{ people survived}]\!])$

(Here and in the following, $0 \leq m, n \leq 600$.) On the other hand, when applied to (10b), (9b) yields:

(12) $\exists m > n$: $g([\![m \text{ people died}]\!]) > g([\![n \text{ people died}]\!])$

which is equivalent to:

(13) $\exists m > n$: $g([\![m \text{ people survived}]\!]) < g([\![n \text{ people survived}]\!])$

It will be clear that (11) and (13) contradict each other.

The third puzzle (I leave the second one for last) was to explain how the following statements can be consistent, and even synonymous:

(14) a. It's good that more than 200 people survived.
b. It's good that fewer than 400 people died.

If co-optation applies, (14a) gives rise to the following inferences:

(15) $\forall m \geq n$:
$g(\llbracket\text{more than } m \text{ people survived}\rrbracket) \geq g(\llbracket\text{more than } n \text{ people survived}\rrbracket)$
$\exists m > n$:
$g(\llbracket\text{more than } m \text{ people survived}\rrbracket) > g(\llbracket\text{more than } n \text{ people survived}\rrbracket)$

As it turns out, co-optation yields exactly the same inferences for (14b). Hence, (14a) and (14b) are equivalent even on the assumption that co-optation applies. Of course, the reason why this outcome is so markedly different from the previous example is that "fewer than" reverses the quantitative ordering on the comparison set associated with "good".

The last remaining puzzle is to explain how (16a) and (16b) manage to be compatible:

(16) a. It's good that 200 people survived.
b. It's bad that 400 people died.

It is instructive to compare this pair to the following examples with non-evaluative gradables:

(17) a. ?Harry is short and tall.
b. Harry is tall for a pygmy but short for a volleyball player.

Whereas it is quite difficult to interpret (17a) as non-contradictory, a plausible construal is readily available for (17b), presumably because the speaker indicates that he is juxtaposing two different measures of height, which he achieves by introducing two different comparison sets. It would be appealing to suppose that a similar shift in perspective distinguishes (16a) from (16b): the same state of affairs is good under one aspect and bad under another. However, if the embedded clauses in (16a) and (16b) are truth-conditionally equivalent, it would seem that the same comparison set is involved in both cases. So how can there be a shift in perspective? The answer, I would like to suggest, is that co-optation makes the difference, for it will induce, in effect, two separate scales on a single set of propositions. (16a) and (16b) evaluate the same state of affairs, but with respect to different backgrounds: in the former case, the statement is relative to a scale of propositions that are ordered from least good to best; in the latter, the same propositions are ordered from least bad to worst.

Having shown how co-optation might help to explain the interpretation of evaluative predicates, it is time to ask ourselves what exactly the status of this principle might be. To begin with, let us consider the possibility that co-optation is somehow hardwired into the semantics of "good", "bad", and related expressions. There are reasons for doubting that this is right. First, as noted at the outset, the construals we've been dealing with aren't always mandatory. In order to take this into account, we would probably have to assume that evaluative predicates are semantically ambiguous between a co-optative and a non-co-optative meaning, which is not an attractive prospect. Secondly, no matter which version of co-optation we adopt, if it was to apply across the board we would predict that evaluative predicates are downward entailing, which doesn't seem correct. To explain, suppose the semantics of "good" is such that "It's good that φ" is true iff $g(\varphi) > s$, where s is a given standard of goodness. Then even the weakest version of co-optation implies that "It's good that ... " is a downward-entailing environment:

(18) If φ is entails ψ, then "It's good that ψ" entails "It's good that φ".

This prediction is dubious, though I should like to note that we must be careful to reject it for the right reasons.

(19) It's good that Edna was found.
$\Rightarrow$ It's good that she was found with a bullet hole in her forehead.

Even if this inference is patently invalid, this doesn't prove that "good" isn't downward entailing. Assuming that the interpretation of "good" is dependent on the comparison set associated with its clausal complement, it is quite likely that the comparison set associated with "Edna was found" will be different from that associated with "Edna was found with a bullet hole in her forehead." Therefore, it is practically inevitable that there will be a shift in perspective when we proceed from the premiss in (19) to the conclusion. Besides, sequences like "It's good that φ, and therefore it's good that ψ" will tend to be infelicitous in any case, simply because "good" is factive. In order to test for the monotonicity properties of "good", it is better to use a non-factive construction and clausal complements that differ from each other merely in quantity:

(20) It would be good if you ate more than 3 apples per week.
$\Rightarrow$ It would be good if you ate more than 5 apples per week.

Now, this looks plausible enough, but then one realises that:

(21) It would be good if you ate more than 3 apples per week.
$\not\Rightarrow$ It would be good if you ate more than 300 apples per week.

That "good" is not downward entailing is confirmed by the observation that it doesn't license negative polarity items:

(22) *It's good that there is any cauliflower left.

If "good" isn't downward entailing, maybe it is upward entailing? This, too, is doubtful:

(23) It would be good if you ate fewer than 300 apples per week.
$\not\Rightarrow$ It would be good if you ate fewer than 3 apples per week.

Hence, "good" appears to be non-monotonic. This conclusion is in line with the following thought experiment. Edna is a great fan of strawberries, but values them most when they come in multiples of 12, and then she doesn't mind if she has 12, 24, 36, etc. Hence, Edna's goodness function for strawberries (or rather, for having strawberries) might look like this:

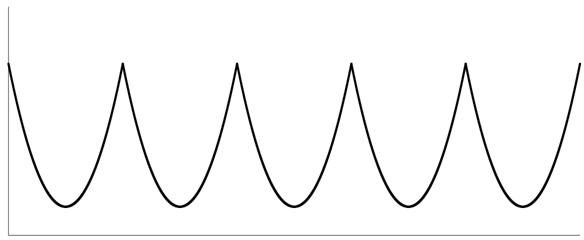

Not that this is a particularly likely scenario, but that is as it may be, as long as we can agree that it is possible in principle. Now Edna says:

(24) It would be good if I had 24 strawberries.

Given Edna's peculiar predilection for multiples of 12, her statement does not entail that it would be even better if she had 25 strawberries, nor does it entail that it would be worse if she had 12.

The moral of the foregoing discussion is that the lexical meaning of "good" doesn't seem to impose any hard constraints on possible goodness functions. However, even if there are no hard constraints, there may well be soft constrains. In fact, I would like to suggest that goodness functions have a default profile:

(25) *Prototypical goodness functions*
Let P = $\varphi_0 \ldots \varphi_k$ be a sequence of propositions aligned according to some quantitative ordering $\preceq$, and $0 < i \leq j \leq k$ (so, possibly, $i = j$ and/or $j = k$). Then a prototypical goodness function for P consists of three subfunctions g, g', and g'', such that:

- dom(g) = $\varphi_0 \ldots \varphi_i$ and g is increasing,
- dom(g') = $\varphi_{i+1} \ldots \varphi_j$ and g' is constant,
- dom(g'') = $\varphi_{j+1} \ldots \varphi_k$ and g'' is decreasing.

Hence, the initial segment of a prototypical goodness function goes up, and then it may level off (if g' is non-empty), and may even take a dive (if g'' is non-empty). A function meeting these specifications could have one of the following contours, for example:

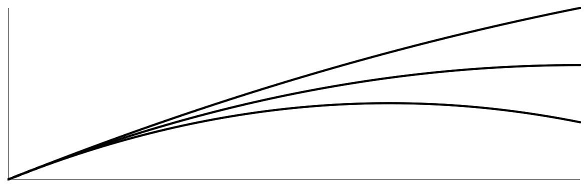

It seems to me that this covers the range of possibilities that readily come to mind when one considers what a goodness function might look like and there isn't much in the way of specific information to go on. In short, this seems like a plausible default to me.

A further assumption that I believe is natural to make is that, by default, if a goodness function has a hanging tail, then the tail will be ignored. For instance, if a speaker says:

(26) It would be good if you ate more than 3 apples per week.

then the hearer is not normally expected to take into account the fact that there is an upper limit to the number of apples that is good for her, even if this is evidently true.

If this story is correct, it follows that, by default, the co-optation assumption holds (though perhaps only within limits), but it doesn't follow that "good" is monotonic. Which would seem to be just the right mix of properties.

To conclude, let me try to say a bit more about the rationale behind co-optation. It has often been remarked that our species has a penchant for establishing connections. If a kangaroo escapes from the local zoo and a few days later another kangaroo does the same, we will immediately wonder whether there might be a connection. Similarly, if a speaker places two events side by side, like this:

(27) Edna fell. Harry pushed her.

hearers will find it hard not to establish a connection. And so on and on.

Co-optation is plausibly seen, I believe, as resulting from the same drive towards coherence. If a speaker associates two orderings with the same set of objects, it is only natural to suppose that the orderings might be related somehow, especially since one of them (the qualitative one) is greatly underdetermined by literal meaning. This explains *why* a connection is made, not *how* it is made. The answer to that question, I would like to suggest, is that co-optation is rooted in world knowledge. Based on regular exposure to quantitative and qualitative scales, we arrive at the notion of a prototypical goodness function, and that is what underlies co-optation.

References

Duchon, D., K. Dunegan, and S. Barton (1989). Framing the problem and making decisions: the facts are not enough. *IEEE transactions on engineering management*: 25–27.

Geurts, B. (2010). Frames and scales. In G. Keren (Ed.), *Perspectives on framing*. Psychology Press.

Levin, I. P. (1987). Associative effects of information framing. *Bulletin of the psychonomic society* 25: 85–86.

Levin, I. P. and G. J. Gaeth (1988). How consumers are affected by the framing of attribute information before and after consuming the product. *The journal of consumer research* 15: 374–378.

Levin, I. P., S. Schnittjer, and S. Thee (1988). Information framing effects in social and personal decisions. *Journal of experimental social psychology* 24: 520–529.

McNeill, B., S. Pauker, H. Sox, and A. Tversky (1982). On the elicitation of preferences for alternative therapies. *New England journal of medicine* 306: 1259–1262.

Seventeenth Amsterdam Colloquium

December 16 – 18, 2009

Pre-proceedings

Organizing Committee:

Maria Aloni, Harald Bastiaanse, Tikitu de Jager,
Peter van Ormondt and Katrin Schulz

ILLC/Department of Philosophy
University of Amsterdam