

WHY ISN'T MY POCKET CALCULATOR A THINKING THING?

Essay, Comments, and Reply

Larry Hauser

William J. Rapaport

Department of Philosophy
Michigan State University
East Lansing, MI 48824

Department of Computer Science
and Center for Cognitive Science
State University of New York at Buffalo
Buffalo, NY 14260

21798lah@msu.bitnet

rapaport@cs.buffalo.edu

January 19, 2005

Abstract

This report consists of three papers: “Why Isn’t My Pocket Calculator a Thinking Thing?”, by Larry Hauser; “Because Mere Calculating Isn’t Thinking” (comments on Hauser’s paper), by William J. Rapaport; and “The Sense of ‘Thinking’,” Hauser’s reply. They were originally presented at the Colloquium on Philosophy of Mind at the American Philosophical Association Central Division meeting in Chicago, 27 April 1991.

Hauser argues that his pocket calculator (Cal) has certain arithmetical abilities: it seems Cal calculates. That calculating is thinking seems equally untendentious. Yet these two claims together provide premises for a seemingly valid syllogism whose conclusion—Cal thinks—most would deny. He considers several ways to avoid this conclusion, and finds them mostly wanting. Either we ourselves can’t be said to think or calculate if *our* calculation-like performances are judged by the standards proposed to rule out Cal; or the standards—e.g., autonomy and self-consciousness—make it impossible to verify whether anything or anyone (save oneself) meets them. While appeals to the intentionality of thought or the unity of minds provide more credible lines of resistance, available accounts of intentionality and mental unity are insufficiently clear and warranted to provide very substantial arguments against Cal’s title to be called a thinking thing. Indeed, considerations favoring granting that title are more formidable than generally appreciated.

Rapaport’s comments suggest that on a strong view of thinking, mere calculating is *not* thinking (and pocket calculators don’t think), but on a weak, but unexciting, sense of thinking, pocket calculators do think. He closes with some observations on the implications of this conclusion.

WHY ISN'T MY POCKET CALCULATOR A THINKING THING?

Larry Hauser

1 INTRODUCTION.

The problem of Other Minds is not just about how one knows that other people are thinking or intelligent beings; it is also about how we know that rocks and roses aren't. From professor to paramecium to paperweight is a slippery slope, and a sense of this is one thing motivating Descartes's notorious views about the mindlessness of lower animals. Thus, he defends his denial that any of them have any mental states at all on the grounds that "there is no reason for believing it of some animals without believing it of all, and many of them such as oysters and sponges are too imperfect for this to be credible" (Descartes 1646: 208). Similarly, I suppose, one reason for denying that any computers have any mental states at all is that there are likewise some of them, e.g., our pocket calculators, "too imperfect for us to believe it of them", or, conversely, it may be feared that if any computing device is rightly credited with any mental abilities at all—even if "only" the ability to add—"then the barriers between mind and machine have been breached and there is no reason to think they won't eventually be removed" (Dretske 1985: 25). Note, here, that the mental abilities we are tempted (if not practically compelled) to attribute to computers are precisely those higher rational abilities, e.g., calculation, which have traditionally been held to distinguish us from, and exalt us above, the lower animals. And it is ironic that so many contemporary arguments would deny mind to machines on the basis of their lack of faculties such as emotion and sense perception that we share with the lower animals, when traditional arguments sought to deny soul and mind to these animals on the basis of their lack of faculties such as mathematical calculation, which, more and more it seems, we are coming to share with computers.

2 CAL.

My inquiry concerns the mental ability—to calculate—of my pocket calculator: call him Cal, for the sake of argument. I suppose most people's intuitions are that their pocket calculators don't think. Yet, we all allow—or speak as if we allow—that they add, subtract, multiply, and divide: in short, calculate. The trouble is that now, given the seemingly obvious thesis that calculation is thinking (indeed, a paradigm case), we have premises for a valid syllogism whose conclusion contradicts "our" original intuitions:

(P1) Calculating is thinking.

(P2) Cal calculates.

(C) \therefore Cal thinks.

Not only do the premises seem true on their face, but obvious. The brunt of my argument will consist in considering various ways in which this argument might be challenged and showing that the costs of each of the strategies proposed is insupportable. Either it turns out that we ourselves can't be said to think or calculate if *our* performances are judged by the standards being proposed to rule out Cal, or else it turns out that the standards are such that it is impossible to verify whether anything or anyone (save, perhaps, oneself) meets them.

The main objections I consider attempt to show that this syllogism commits the fallacy of four terms on the grounds that ‘calculation’ is only equivocally predicable of Cal, that Cal doesn’t *really* calculate because Cal’s performances lack features essential for genuine cognition or calculation. These features generally fall under the headings of four traditional “marks of the mental”: consciousness, autonomy, intentionality, and unity. Of these, the last two underwrite the most serious challenges to Cal’s claim to be calculating, and hence a thinking thing; yet, in the end, I urge that neither of these objections is all that compelling either, nowhere near so compelling as our original premises. What is amiss, I urge, is “our intuition” that our pocket calculators don’t think: I suggest that this intuition is as theory laden (and perhaps theology laden) as Descartes’s “intuition” that oysters don’t have souls.

3 CONSCIOUSNESS.

The argument from consciousness holds that the essence of thinking is its subjectivity: there must be something that it’s like to be a pocket calculator *for* the calculator, or else it’s not really calculating. The trouble with this objection is that it can’t be *substantiated* just how far (beyond myself) this mysterious “inner light” of consciousness extends. This “other minds” reply does not, as Searle (1980: 422) jeers, “feign anesthesia”: it only requires critics of AI to consistently apply the criterion they propose to disqualify computers’ claims to think. What the Other Minds Reply says is that *if* consciousness *were* our basis for deciding whether any intelligent-seeming thing were really a thinking subject, then one *should* have skeptical doubts about other minds. So, if we don’t, and shouldn’t, seriously entertain such doubts, this seems to show that we don’t (or shouldn’t) appeal to consciousness to decide what is and isn’t thinking.

The general argumentative strategy is that no proposed criterion of thinking is acceptable if its application leads to the conclusion that people don’t think, or that we have no way of telling this, in cases where we think we know well enough that they do. The demand that criteria proposed be applied consistently to human and machine, and not selectively to machines, is unexceptionable: otherwise one begs the question.

4 AUTONOMY.

Autonomy objections, like the Objection from Consciousness, touch some deep chord; and such objections—that computers lack freedom or wills of their own, that they “can only do what we tell them”—are among the most frequently heard arguments against artificial intelligence. Autonomy, as a criterion for distinguishing genuine from apparent cognition, faces problems akin to those that arise for consciousness. If one appeals to introspection—if our awareness of and basis for attributing free will or autonomy is supposed to be phenomenological—then autonomy objections inherit all the problems of the Objection from Consciousness: it would be impossible to know (or even have justified belief) that anything or anyone (save oneself) really is a thinking subject of mental states. On the other hand, if we reject the phenomenological criterion of autonomy, as I suspect we should, the Autonomy Objection becomes even less supportable. *With* consciousness of autonomy as the criterion of autonomy, it seems we can never know that anyone *else* has it. Without consciousness, it seems we can’t even know that we ourselves do. Note that the burden of the Autonomy Objector here is not just to show that there *are* free acts in some strong libertarian sense—though this would be onerous enough—but to show that certain acts *of ours* are free in this sense and no acts of computers are or (presumably) ever could be. I have no idea how—without appeal to one’s

introspective sense of freedom as evidence of actual freedom—one might propose to discharge this burden.

If a conscious or introspective sense of freedom is supposed to evidence freedom, the Autonomy Objection can bear no more weight than the appeal to consciousness it presupposes. Yet without such appeal, our own claims to be thinking could be no better grounded (and might be worse) than libertarian metaphysical doctrine. Yet my beliefs that I have beliefs, desires, and such—even my belief that you do—seem nowise so dubious as that.

5 INTENTIONALITY.

A third line of objection to the claim that computers generally—and Cal, in particular—really think, appeals to the intentionality of mental states. The Intentionality Objection, as Searle (1980) puts it, is that the symbols or information that computers process are only symbolic *of* something, or information *about* anything, *to us*, that they are not of or about anything to the computer. So stated, the Intentionality Objection threatens to collapse into the Consciousness Objection—if the difference between my calculation and Cal’s “calculation” is just supposed to be that there is something that it’s like for me to calculate that $2 + 9$ is 11, but nothing that it’s like for Cal. Note how closely what Searle says about the computer’s relations to the symbols and information it processes—that they’re only symbolic and informative *for us*, not for the computer—echoes formulations such as Nagel’s (1974) or Sartre’s (1956) about what it is to be conscious. To go on from this to deny the possibility of any behavioral or public tests of intentionality (to deny these inward states any outward criteria), as Searle seems to do, really does seem to make the Intentionality Objection into a (species of) Consciousness Objection. To avoid this collapse, one has to indicate factors (besides consciousness) that distinguish unthinking “syntactic” manipulations from contentful thought, and the factors in question must be public or observable—the sorts of things Wittgenstein referred to as “outward criteria” (Wittgenstein 1958, §580). For this reason, Dretske’s (1985) version of the Intentionality Objection has much to recommend it over Searle’s.

What Dretske takes to be the missing ingredient—what we have that computers lack—are causal connections between the signs and the things they signify. Put crudely, the difference between my contentful belief *that dogs are animals* and a computer’s “representation” of this same information—say by storing a Prolog clause [that says] `animal(X) :- dog(X)` in RAM—is that my representation came to be, or could be, elicited by the actual presence—the actual sights and sounds—of dogs. It is these perceptually mediated connections between my tokenings of the English word ‘dog’ and actual dogs that makes that word signify those animals for me, and it is for want of such connections that computer representations—tokenings, say, of the Prolog atom `dog`—lack such signification for the computer. Here, Dretske sets out an Intentionality Objection that stands or falls independently—as Searle’s formulations, do not—of appeals to consciousness.

Yet despite the superiority of Dretske’s formulation (in providing a positive account of signification), it has two serious shortcomings. First, it has less force than Dretske seems to think, even in the most favorable cases, i.e., of signs for perceptible things such as dogs or (perhaps most favorably of all) perceptible qualities such as color or pitch. Second, not all words or signs are as rich in sensory associations as ‘red’ or ‘dog’. Where signs, such as numerals, represent abstractions, such as numbers, it seems less plausible to think significance requires causal links to the things signified or any very robust causal-perceptual links with anything.

With regard to the first, most favorable, cases, Dretske's account of reference threatens the conclusion that perceptually deficient humans are meaning-deprived also. Presumably (on this view) someone blind from birth cannot signify anything by color words, nor the congenitally deaf by words like 'pitch' and 'music'. Yet I believe there are good reasons (see, e.g., Landau & Gleitman 1985) to hold that such persons can use words to refer to such things despite their lack of perceptual access to them. Someone blind from birth, e.g., could know *that fire engines are red, that red is a color, and even that 'color' refers, roughly, to those qualities of things, deriving from their reflective properties, which the sighted see, but I don't, and which stand in a similar relation to vision as tone to hearing*. More dramatically, on Dretske's view, how are we to avoid the absurd consequence that most of the words in Helen Keller's books, though signifying things to us, signified nothing (or at least much less) to Helen Keller? Even for words with very considerable ostensible or sensuous content, the ability to apply these words to the world on the basis of this content seems less crucial to signification than Dretske's views suggest. Perhaps some causal-perceptual associations for some terms in one's lexicon are prerequisite for signifying anything by any term at all, but Helen Keller's case seems to suggest that these connections needn't be so direct or robust as Dretske has to require to make his case against (present-day) machine understanding.

But whatever the upshot of these reflections on the most favorable cases for Dretske—terms ordinarily rich in ostensive or sensuous associations—plainly, cases such as the numerals are less favorable. Whatever weight Dretske's appeal to the perceptual disabilities of (present-day) digital computers has against their claims to mean anything by words such as 'dog' and 'red', this argument will be impotent to establish the more general claim Dretske wants—that none of the symbols (present-day) computers process signify anything to these computers—if Cal's understanding of the numerals he processes, e.g., is immune to such objections. How is Dretske's causal-perceptual account of reference supposed to work here, where the referents, numbers, don't seem capable either of being perceived or entering into causal relations?

At this point, the only plausible move for a causal-perceptual theory of reference seems something like this: in order to mean numbers by numerals, one must be able to apply numerals to items and events in the world, e.g., in counting; so, in order to mean two by '2', say, one must be capable of reliably tokening '2' when presented various pairs of objects or events. Yet even if this is correct, and causal-perceptual links are in some sense required for reference even to numbers, the senses of 'perceptual' and 'in the world' here cannot be very robust. Someone in a state of total (external) sensory deprivation might still count their breaths or even how many times they performed a carry in doing a bit of mental addition, and if this is all that's required, it's clear that it's not enough to rule out Cal. If all that's required "of any aspiring symbol manipulator is, in effect, that *some* of its symbols be actual signs of the conditions they signify, that there be some symbol-to-world correlations that confer on these symbols an intrinsic meaning" (Dretske 1985: 29), then so long as Cal can count key presses and iterations of loops, this would seem causal-perceptual linkage enough to support his claim to mean numbers by his numeric tokens.

If this is right, it means the Intentionality Objection cannot be sustained across the board, with regard to all the symbols computers process. Perceptually impoverished as Cal is, it seems Cal has enough reality contact to support his claim to mean two by [the binary numeral] '10'.

6 UNITY.

Since Plato, at least, unity has been advanced as a distinguishing attribute of minds: The idea is that minds are not composed of parts (as bodies are), but are rather indissoluble units, a claim that not only underwrites perhaps the most important traditional argument for the immortality of the soul, but which has been continually in the background of our whole discussion to this point. Each of the various objections we considered makes some tacit appeal to unity: each claims that disconnected from some further mental abilities or events (consciousness or autonomy or intentionality) Cal's seeming calculation is not really calculation at all. What distinguishes these abilities as "marks of the mental" deserving special treatment—as compared to, say, the ability to enjoy the taste of strawberries—is just their greater intuitive and traditional claims to being ineliminable aspects of all mental activities or states.

Having been unsuccessful in sustaining objections to Cal's claim to calculate on the basis of any of these specific abilities, we turn now to a nonspecific form of this whole general line of objection. With regard to Cal's calculation, a proponent of the Unity Objection might respond that if calculation is the *only* cognition-like thing Cal can do, then it's not thinking and perhaps not really even calculation. The Unity Objection, rather than claiming that some specific mental ability is necessary for thought or calculation, claims what is essential to something's thinking, or even calculating, is having *enough* interconnected mental abilities of various sorts. Even if we stop short of the strong Cartesian demand that thought must be "a universal instrument which can serve for all kinds of situations" (Descartes 1637: 140), perhaps we can at least require that would-be thought processes should be *flexible* instruments that can serve for *various* situations. This seems sufficient to exclude Cal without excluding me or you or Helen Keller.

On the other hand, while this *may* suffice to exclude hard-wired special purpose devices like Cal from the ranks of the thinking, it does not seem so effective against the claims of programmable machines, such as my laptop computer. Indeed, much of the deep philosophical interest of AI derives from the fact that programmable digital computers are in fact flexible and even—in the sense that "they can mimic any discrete state machine" (Turing 1950: 441)—universal instruments.

This presents us with a conundrum: Suppose my laptop computer—call her Sparky—were programmed to emulate Cal. Suppose Sparky computes the same arithmetical functions, by the same procedures, as Cal. Now it seems odd to say that Sparky calculates, but Cal doesn't, just because Sparky has other abilities (or at least can be programmed to have other abilities). If both compute the same functions using the same algorithms, aren't they—in the sense relevant to cognitive attribution—doing the same thing? Perhaps the Unity Objection, for all its traditional and intuitive warrant, is misguided. As Dretske remarks, "We don't, after all, deny someone the capacity to love because they can't do differential calculus. Why deny the computer the ability to solve problems or understand stories because it doesn't feel love, experience nausea, or suffer indigestion?" (Dretske 1985: 24).

What the Unity Objection seems to require—and offer no prospect that I can see of providing—is some account not only of *how many* and *which* other mental abilities a thing must have in order to calculate (or think), but *why*. If Cal follows the same addition procedure as Sparky, and Sparky the same addition procedure as I, then it seems to me that Cal adds if I do; and, *when we do*, 'calculation' is predicable of both of us in exactly the same sense, regardless of whatever further mental abilities of mine Sparky lacks or whatever further capacities of Sparky's are absent in Cal. Nor is it even essential that the procedures Cal and Sparky follow should emulate those I,

or people generally, follow. This is not Searle’s Brain Simulator reply but the reverse—the CPU Simulator reply—that it’s enough that the procedure that Cal follows be one that I *could* follow (e.g., by hand-simulating Cal’s processing) and that in doing this *I* would be calculating.

7 CONCLUSION.

What the preceding arguments show—I take it—is that none of the four traditional marks of the mental considered provide a supportable basis for denying that Cal calculates in the same sense as you or I; i.e., I have sought to show that our initial syllogism does not commit the fallacy of four terms by equivocating on ‘calculates’, its middle. I will conclude by remarking why the argument—at least as I intend it, and on its least tendentious reading—doesn’t equivocate on its major, ‘thinks’, either. Ordinarily ‘think’ is a generic term for any of several different mental activities or states. According to Descartes, a thing that thinks is “a thing which doubts, understands, affirms, denies, is willing, is unwilling, and also imagines and has sensory perceptions” (1642: 19). Similarly, my dictionary (*Webster’s New Collegiate*), under ‘think’, mentions conceive, judge, consider, surmise, expect, determine, resolve, reason, intend, purpose, reflect, infer, opine, and decide. In this ordinary generic sense of the term, I take it, it’s undeniable that calculating is thinking, and—if my arguments are sound—that my pocket calculator calculates and consequently thinks.

Perhaps some special sense of “thinking” can be made out for which calculating is not sufficient—perhaps some sense in which it’s not sufficient to doubt *or* understand *or* will, etc., but in which it’s necessary to (be able to) doubt *and* understand *and* will, etc. (as Descartes surely intended). Perhaps there is some sense in which “thinking” requires such unity, or universality of mental capacity—or alternatively some other traditional (or perhaps some non-traditional) mark(s) of the mental. At any rate—whether or not such a sense of ‘thought’ can be made out—I have only claimed that Cal thinks in the ordinary generic sense of being a subject of at least one kind of contentful or mental state, not that he is a unified, or conscious, or autonomous self or soul or thinker in some special proprietary philosophical sense. I leave it to the opponent of AI to clarify what this sense is and to make out the case, if it can be made, against Cal’s thinking in *this* sense.

REFERENCES

1. Descartes, René (1637), *Discourse on Method*, trans. in J. Cottingham, R. Stoothoff, & D. Murdoch, *The Philosophical Writings of Descartes*, Vol. 1 (Cambridge, Eng.: Cambridge University Press, 1985): 131–141.
2. Descartes, René (1642), *Meditations on First Philosophy*, trans. in J. Cottingham, R. Stoothoff, & D. Murdoch, *The Philosophical Writings of Descartes*, Vol. 2 (Cambridge, Eng.: Cambridge University Press, 1984): 16–23
3. Descartes, René (1646), “Letter to the Marquess of Newcastle, 23 November 1646,” trans. in A. Kenny, *Descartes’s Philosophical Letters* (Oxford: Clarendon Press, 1970): 205–208.
4. Dretske, Fred (1985), “Machines and the Mental,” *Proceedings and Addresses of the American Philosophical Association* 59: 23–33.

5. Landau, Barbara, & Gleitman, Lila (1985), *Language and Experience: Evidence from the Blind Child* (Cambridge, MA: Harvard University Press).
6. Nagel, Thomas (1974), "What Is It Like to Be a Bat?", *Philosophical Review* 83: 435–450.
7. Sartre, Jean-Paul (1956), *Being and Nothingness*, trans. by H. Barnes (Secaucus, NJ: Citadel Press).
8. Searle, John R. (1980), "Minds, Brains, and Programs," *Behavioral and Brain Sciences* 3: 417–457.
9. Turing, Alan M. (1950), "Computing Machinery and Intelligence," *Mind* 59: 433–460; reprinted in A. R. Anderson (ed.), *Minds and Machines* (Englewood Cliffs, NJ: Prentice-Hall, 1964): 4–30.
10. Wittgenstein, Ludwig (1958), *Philosophical Investigations*, trans. by G. E. M. Anscombe (Oxford: Basil Blackwell).

BECAUSE MERE CALCULATING ISN'T THINKING:

Comments on Hauser's "Why Isn't My Pocket Calculator a Thinking Thing?"

William J. Rapaport

1 DO POCKET CALCULATORS THINK?—HAUSER'S MASTER ARGUMENT.

The first time I read Larry Hauser's essay, "Why Isn't My Pocket Calculator a Thinking Thing?" (Hauser 1991), I thought he was obviously wrong, but then my colleagues convinced me that he was right. So, in these comments, I want to say something about both of these intuitions.

Hauser argues that:

- (P1) Calculating is thinking.
- (P2) Cal [a pocket calculator] calculates.
- (C) \therefore Cal thinks.

He defends this perhaps counterintuitive conclusion against four objections. Now, both Hauser and I would agree, I think, that a *computer that is executing a suitable AI program* that includes modules for calculating and that passes the Turing Test can be said to think. That's controversial, of course, but rather less so than Hauser's present claim about a very much simpler machine, a claim at which even many researchers in the so-called strong-AI tradition might balk. I will argue that two of the objections are successful against Hauser's argument on one sense of thinking, but that on a weaker sense, Hauser's argument is sound but unastounding.

We must first clarify two terms: What do 'thinking' and 'calculating' mean? 'Thinking' presumably means something like "exhibiting mental abilities". This could be a wide spectrum from exhibiting all or most mental abilities to exhibiting only some or even one. At the minimal end of this spectrum, where 'thinking' means exhibiting at least one mental ability—the sense in which someone who is adding, say, can be said to be thinking—Hauser's argument is sound but not the controversial claim that it appears to be at first sight (or at most it's a controversial claim about how to use the word 'thinking'). At the maximal end of the spectrum, where 'thinking' means exhibiting all or most mental abilities, thinking is roughly equivalent to passing the Turing Test—this is the sense in which 'thinking' is a cover term for all mental processes (i.e., for intelligence in general). On this view, (P1) and (C) are false.

Next, what does 'calculating' mean for Hauser? Two interpretations come to mind:

- (I1) Calculating is doing arithmetic (following or executing algorithms for $+$, $-$, \times , \div).
- (I2) Calculating is computing (in the sense of the Church–Turing Thesis).

To see which of these Hauser might mean by 'calculate', consider (P2), the premise that Cal calculates. Hauser says:

My pocket calculator (Cal) has certain *arithmetical* abilities: it seems Cal calculates. (Abstract; my italics.)

... we all allow—or speak as if we allow—that they [pocket calculators] add, subtract, multiply, and divide: in short, calculate. (p. 2.)

These surely favor interpretation (I1), that to calculate is to do arithmetic.

But interpretation (I1), that to calculate is to do arithmetic, may not be quite what Hauser has in mind. Rather, interpretation (I2), that to calculate is to compute, may be what he has in mind, since computation is often linked to thinking in discussions of AI. Consider these claims:

That calculating is thinking seems equally untendentious. (Abstract.)

... the seemingly obvious thesis that calculation is thinking (indeed, a paradigm case) ... (p. 2; this quotation comes almost immediately after the other quotation from p. 2, cited above.)

I would have thought that something more complex, like natural-language understanding, would be a “paradigm case” of thinking.

In any event, what might (P1), that calculating is thinking, mean on the two interpretations of ‘calculating’? Note, first, that ‘calculating is thinking’ is ambiguous: It could mean that whatever calculates thinks, or (perhaps) that whatever thinks calculates (i.e., that calculating is what thinking is), or that calculating is identical with thinking. So we have five cases:

(P1.I1.a) *Doing arithmetic is a kind of thinking.* On the *maximal* view of thinking, a subject of the Turing Test that can do arithmetic does not *ipso facto* pass the Turing Test, nor does a subject that cannot do arithmetic *ipso facto* fail. So, by the Turing Test criterion, doing arithmetic is neither necessary nor sufficient for thinking. Thus, on the (I1)-interpretation of calculating as doing arithmetic, the allegedly “untendentious” premise (P1), that calculating is thinking, is indeed tendentious. However, on the *minimal* view of thinking—the view on which if someone were calculating, you could say truly of them that they were thinking—(P1) appears to be acceptable, if unexciting.

(P1.I1.b) *To think is to do arithmetic.* This is clearly false, but I put it here for the sake of completeness.

(P1.I2.a) *Computing is a kind of thinking.* By the fundamental assumption (or working hypothesis, if you prefer) of computational cognitive science, *all thinking is computing but not vice versa* (cf. Pylyshyn 1985, Rapaport 1990b). So, “if something computes, then it thinks” is false. So (P1.I2.a) is false. Thus, on the (I2)-interpretation of calculating as computing, (P1) is still tendentious.

(P1.I2.b) *To think is to compute.* The converse of (P1.I2.a), which, as noted above, is the fundamental assumption of computational cognitive science, is, *I think*, true, but it is clearly tendentious.

(P1.=) *Calculating = thinking.* This is unacceptable on either interpretation, since neither the pair (P1.I1.a)–(P1.I1.b) nor (P1.I2.a)–(P1.I2.b) is untendentiously acceptable.

So, for Hauser’s argument to be valid, premise (P1), that calculating is thinking, must be understood as (P1.I1.a) on the minimal reading of ‘thinking’. Let us now turn to the objections.

2 OBJECTION 1: THE ARGUMENT FROM CONSCIOUSNESS.

... the essence of thinking is its subjectivity: there must be something that it's like to be a pocket calculator for the calculator, or else it's not really calculating [*sic*]. (p. 3)

I take this objection to attack (P1), that calculating is thinking, on interpretation (P1.I1.a), that calculating is doing arithmetic, as follows (Hauser seems to interpret the force of this objection a bit differently; see section 6, below):

(O1) Thinking requires subjectivity. (“the essence of thinking is its subjectivity”)

(O1A) Calculating does not require subjectivity. (Missing assumption.)

\neg (P1.I1.a) \therefore Calculating isn't thinking.

Hauser's “Other Minds Reply” to (O1) is that if (O1) were true, we'd never know whether *anything* were a thinking thing and that we should not entertain such skeptical doubts about other minds. So (O1) should be rejected. I can agree with that.

I might add the following observations: As to (O1) itself, I grant that subjectivity is an important issue (cf. Rapaport 1990a, Galbraith & Rapaport 1991). For instance, suppose that I use a calculator to add two numbers while simultaneously adding the two numbers in my head; in such a case, I am executing two distinct processes, at least one of which (the processing in my head) exhibits subjectivity. If I never learned how to add, but were given instructions on how to use the calculator, then by using it I would be adding. Yet I could say that I myself did not know how to add.¹ That would not contradict the fact that, by pushing calculator buttons, I am adding. But is subjectivity exhibited in this case? Or would it be better to say merely that there is an adding process going on? The question raised here—echoing Bertand Russell's objection² to Descartes's *cogito ergo sum*—is: “Where does *I* come from?”³

But does thinking *require* subjectivity? On one interpretation, (O1) is the absent-qualia problem. Since I hold that qualia *can* be absent—roughly, because I take qualia to be implementation-dependent side-effects, and a particular implementation might not provide any such side-effects—I discount the strength of (O1).

Premise (O1A) is the missing assumption needed to yield \neg (P1.I1.a). But there is no evidence to support it: How would one know whether calculating *requires* subjectivity? (How would one know what it is like to be Cal?) So (O1A) is dubious (though probably true!).

3 OBJECTION 2: THE ARGUMENT FROM AUTONOMY.

... computers lack freedom or wills of their own, ... they “can only do what we tell them” ... (p. 3.)

I take it that this is also an objection to (P1.I1.a), as follows:

¹This, no doubt, has ramifications for contemporary mathematics education!

²Cf. Russell 1950, as cited in Kenny 1968: 58-59.

³With apologies to Mary Galbraith for usurping her phrase.

(O2) The sort of calculating that computers (and, presumably, calculators) do is not autonomous. (They “can only do what we tell them”.)

(O2A) Thinking requires autonomy. (Missing assumption.)

¬ (P1.I1.a) ∴ Calculating isn’t thinking.

Hauser’s reply to (O2) is that, first, if our evidence for autonomy is introspective, then there is no way to tell if *other* minds are autonomous (which is bad), and, second, if our evidence for autonomy is *not* introspective, then there is no way to tell if *we* are autonomous (which is worse). I agree, although I have some doubts about the relevance of “autonomy” and “free will” issues to the ability of computers to “only do what we tell them”. If “[they] can only do what we tell them” means that they are not “free” to do things “on their own”—that they are “limited” or “determined” in their abilities by their physical or (hardwired) algorithmic makeup, then one might reply: So are we! Thus, either we both can think, or neither of us can!

On another interpretation, if “[they] can only do what we tell them” means that their behavior is entirely predictable, then I would merely note that pocket calculators can do things that we don’t expect (just as chess computers, natural-language-understanding computers, etc., do): There are lots of undocumented things one can do with a pocket calculator—things not explicitly or intentionally programmed in but that are logical consequences (side effects?) of the program. For example, on many pocket calculators, you can raise a number to a power by inputting the number, pushing the ‘×’ button, and repeatedly pushing the ‘=’ button; this is often undocumented. So, there is a sense in which Cal can do more than it is told.

4 OBJECTION 3: THE ARGUMENT FROM INTENTIONALITY.

... the symbols or information that computers process are only symbolic *of* something, or information *about* anything, *to us*; ... they are not of or about anything to the computer. (p. 4.)

Again, this appears to be an objection to (P1.I1.a):

(O3) The sort of calculating (namely, symbol processing) that computers (and, presumably, calculators) do lacks intentionality. (I.e., Cal’s symbols are only meaningful to *us*.)

(O3A) Thinking requires intentionality. (Missing assumption.)

¬ (P1.I1.a) ∴ Calculating isn’t thinking.

This objection implicitly and explicitly echoes the arguments of Searle’s Chinese Room (1980) and Dretske’s “Machines and the Mental” (1985).

Hauser’s version of Dretske’s version of (O3) is that the calculator lacks, but we have, “causal connections between the signs and the things they signify” (p. 4). Hauser argues against (O3) that causal connections aren’t essential—witness blind people’s use of color terms. I am sympathetic to this (for the details, see Rapaport 1988, esp. sect. 3.3.), but the issue of a blind person’s use of color terms, it seems to me, is a complex issue that cannot adequately be dealt

with here (cf. Crimmins 1989 for further discussion, and cf. Maloney 1987). In any event, even Helen Keller had *some* causal connections, via touch. Granted, her use of words would have been different from ours, but so are mine from yours, yet we can communicate (cf. Rapaport 1988, sects. 3.5, 4.2).

But Hauser goes on to observe that

... so long as Cal can count key presses and iterations of loops, this would seem causal-perceptual linkage enough to support his claim to mean numbers by his numeric tokens. (p. 5.)

Here, I think, Hauser's argument can be strengthened. For one thing, Cal *cannot* count key presses or loop iterations; Cal, after all, is merely a pocket calculator. But Cal does have the relevant intentionality: Cal can link numerals with numbers: The output-numerals are on the display, and the input-numerals are the buttons; numbers are internal data structures implemented in its wiring. (I owe this point to Stuart C. Shapiro. Arguably, these are numerals, too, but I won't press that point, since I don't think it's relevant—see Rapaport 1988.) And Cal can link operation signs (also the buttons) with operations (internal actions, also hard-wired in). Perhaps Cal's symbols mean something different for it than for us; so be it. But as long as Cal performs the right algorithms on the right input, Cal's calculating is relevantly intentional.

Now, perhaps what Cal lacks is a more complete theory of arithmetic and an understanding of what it's doing. These are part of the maximal notion of thinking. On this view, Cal's calculating does arguably lack the relevant intentionality, for not only does Cal lack the ability to count key presses and loop iterations, it cannot link these acts of counting to each other or to arbitrary, independent arithmetic problems. Neither can it reasonably be said to "know" what adding is, nor to "know" what the number 1 is nor what the numeral '1' is, nor their relationship. Thus, it is not able to provide its own input and interpret its own output. In short, it lacks a full theory of arithmetic, precisely what would be needed to "promote" mere calculating to maximal thinking. On this view, (P1) is successfully attacked by (O3), for the case of Cal. I would argue, and I'm sure that Hauser would agree, that a suitably AI-programmed *computer* with a full theory of arithmetic and the self-sufficiency to provide its own input (to set its own problems) and to interpret its own output would be immune to (O3). This is debatable, of course, but now is not the time to debate it (for details, see Rapaport 1988).

So, on the minimal view of thinking, Cal *has* intentionality and therefore *can* think and calculate; but on the maximal view of thinking, Cal can't think.

5 OBJECTION 4: THE ARGUMENT FROM UNITY.

... minds are not composed of parts (as bodies are), but are rather indissoluble units
(p. 6.)

This attacks (P1.I1.a) more or less as follows:

(O4) Thinking things are indissoluble units.

(O4A) Calculating is dissoluble.

\neg (P1.I1.a) \therefore Calculating is not thinking.

Objection (O4) strikes me as a rather vague, outdated, and irrelevant mark of the mental. Moreover, if the mental *is* computational, then it *is* “dissoluble” in the sense that any computation can be expressed in terms of—i.e., can be analyzed or “dissolved” into—basic operations and complex operations formed by such “control structures” as sequencing, selection, and repetition (as expressed in the Boehm-Jacopini and Kleene Normal Form theorems about computation as analyzable into basic operations plus control structures; see, e.g., Marcotty & Ledgard 1986). On this view, (O4) is false, and (O4A) would, of course, be acceptable. (I am not at all sure what (O4) might otherwise mean.)

But Hauser offers a stronger version of this objection:

... if calculation is the *only* cognition-like thing that Cal can do, then it’s not thinking
[W]hat is essential to something’s thinking ... is having *enough* interconnected mental abilities of various sorts. (p. 6.)

A stronger form of the Argument from Unity, then, appears to be the following:

(O4S) Thinking involves interconnected mental abilities of various sorts.

(O4SA) Calculating does not involve (other) mental abilities of various sorts.

\neg (P1.I1.a) \therefore Calculating isn’t thinking.

This is much better, though I’d put (O4S) slightly differently: The more mental things a computer can do, the more it can (be said to) think, i.e., the more maximal its thinking is. The sort of thinking involved in (O4S) appears to be maximal thinking.

But then Hauser undercuts his own argument!:

... *while this may suffice to exclude hard-wired special purpose devices like Cal from the ranks of the thinking*, it does not seem so effective against the claims of programmable machines, such as my laptop computer. (p. 6; my italics.)

But his original argument *was* about *Cal*, *not* about a more complex device.

Even worse, now Hauser seems to have changed the subject matter of the argument: Does he really think that (O4S) is not effective against an arbitrary general-purpose computer? Can Hauser’s argument be made to go through for an arbitrary PC? I’ll return to that later (section 7).

So (O4S), by Hauser’s own admission, is fatal to (P1) on the maximal-thinking interpretation. But why should we believe (O4S)? If an arbitrary PC (“Sparky”) emulates Cal, does Sparky calculate? Clearly, I think, the answer is ‘yes’, at least insofar as Cal also calculates (for details of an argument to this effect, aimed at showing that Searle in the Chinese Room does understand Chinese, see Rapaport 1990a). But does Sparky *think*? And does it think just because it can do other things? Even on the maximal interpretation of thinking, I would say ‘no’, because, first, the other things that Sparky can do (e.g., word processing) *aren’t mental abilities*. One might be willing to say that calculating is a mental ability. But is word-processing a mental ability? Second, the other things that Sparky can do *aren’t* linked to the calculating program; Sparky’s is not a unified mind: Sparky doesn’t have the right *kind* of “interconnected mental abilities”.

It’s perhaps worth noting that the notion of “various sorts” of mental abilities is not at all clear. On a Fodorian modular view of mentality, the “various sorts” might be the ability to do other things besides merely calculating. Arguably, however, these other things are not *linked*—they are modular (Fodor 1983) or cognitively impenetrable (Pylyshyn 1985). But possibly all mental abilities are ultimately the same; e.g., if thinking is computing, then all different kinds of thinking are ultimately expressible in terms of the primitives and constructs of, say, a Turing Machine. It might even be the case that thinking is that subset of computing that is capable of expressing arithmetic; if so, then Cal would be able to exhibit all mental abilities (including even understanding natural language, by some *very* long and complex sequence of ordinary arithmetic operations with data suitably encoded by numerals)!

6 IS (P1) OR (P2) AT FAULT?

Up till now, I have taken it as obvious that the Four Objections to the argument that Cal thinks have been objections to the truth of premise (P1), that calculating is thinking. Yet at the end of Hauser’s discussion of the Four Objections, he tells us that

... none of the four traditional marks of the mental ... provide a supportable basis for denying that *Cal calculates* in the same sense as you or I ... (p. 7, my italics.)

But, surely it is clear that Cal calculates and that none of (O1)–(O4) attack that. On what basis could one possibly deny that Cal was calculating? On the basis that it is not *Cal* that calculates, but a person who *uses* Cal to calculate? One might argue that an abacus doesn’t calculate, but that only a system consisting of a human *using* an abacus calculates (cf. the “Systems Reply” to Searle’s Chinese Room Argument). But if anything that executes an algorithm for calculating can be said to calculate (and I think that that is the case), then if I followed Cal’s algorithms, I would be calculating, too. So, Cal calculates.

In any event, should the Four Objections have been taken as attacking (P2), not (P1)? Indeed, at the very beginning of Hauser’s paper, he told us that

The main objections I consider attempt to show ... that *Cal doesn’t really calculate* because Cal’s performances lack features essential for general cognition or calculation. (p. 3, my italics.)

Can the Argument from Consciousness be taken as an argument against (P2)? But it only speaks of thinking and subjectivity, and in the one place where it does mention calculating (“there must be something that it’s like to be a pocket calculator for the calculator, or else it’s not really calculating”), surely Hauser meant to say that there must be something that it’s like to be a pocket calculator for the calculator, or else it’s not really *thinking*. So (O1) should not be taken as an argument against (P2).

What about the Argument from Autonomy? Since it only talks about cognition and thinking, not calculating, it, too, is best taken as an attack on (P1), not (P2).

The Argument from Intentionality explicitly states that it is an objection to the claim that Cal can *think* (C), so it can be construed either as an attack on (P1) or on (P2).

Finally, the Argument from Unity explicitly states that it is an objection to the claim that calculating is thinking (P1).

Thus, Hauser’s argument construed as attacking (P2) (and as taking (P1) as trivial) is weak. His argument construed as attacking (P1) (and as taking (P2) as trivial)—which is how I have construed it—is more interesting. But it only works for the unexciting case of minimal thinking and fails for the more interesting case of maximal thinking, by the Arguments from Intentionality and the strong version of Unity.

Perhaps minimal thinking leads inexorably, by a slippery slope, to maximal thinking. Hauser observes that

... I have only claimed that Cal thinks in the ordinary generic sense of being a subject of *at least one* kind of contentful or mental state, not that he is a unified, or conscious, or autonomous self or soul or thinker ... (p. 7.)

This, then, is the minimal sense of thinking. But for that claim to be the one that “most would deny” (Abstract) on the grounds that it “contradicts ‘our’ original intuitions” (p. 2), it would have to be the very much stronger claim that Cal thinks in the same sense that we do—the maximal sense—not in some watered-down sense. So, at best, what Hauser’s original argument shows is that Cal does one of the things that counts as thinking. But it does not follow from that that Cal thinks in the maximal sense. Hauser’s argument from Descartes’s and *Webster’s New Collegiate Dictionary’s*⁴ definitions (p. 11) only works if those *conjunctive* definitions are replaced by disjunctive ones! These definitions define ‘thinking’ by a conjunction of kinds of thinking (e.g., doubting, understanding, etc., for Descartes, and conceiving, judging, etc., for the dictionary). Thus, to think must be to exhibit *all* kinds of thinking, not just one. This brings us back to the one remaining open question.

⁴Note that *Webster’s New Collegiate Dictionary* takes a “descriptivist” stance towards definitions, hence ought not to bear any weight in philosophical argument.

7 DOES SPARKY THINK?

Does an argument similar to Hauser’s original one show that a PC (or, in general, a Turing Machine) thinks? Consider:

(P1.I2.a) Computation is thinking.

(P2’) Sparky (a PC) computes (more precisely, Sparky can (be programmed to) do anything that’s computable).

(C’) \therefore Sparky thinks.

Premise (P2’) is clearly true, by Church’s Thesis. Premise (P1.I2.a), on the other hand, is much too strong, as I argued earlier (in section 1). No one, to my knowledge, has claimed (P1.I2.a); many have claimed its converse. But, for the sake of argument, let’s apply the Four Objections to (P1.I2.a).

Could one deny it by (O1), that there is nothing it is like to be Sparky computing? No; Hauser’s and my objections to (O1) hold.

Could one deny it by (O2), that Sparky can only do what we program it to do? No; Hauser’s objection to (O2) holds, as does mine: (O2) is just false if it means that programmers can’t be surprised by their own programs.

Could one deny it by (O3), that Sparky lacks intentionality? I would say ‘no’ if Sparky is running a (suitable) AI program, but I would say ‘yes’ if Sparky runs any other kind of program: Intentionality is not required for arbitrary computation.

Finally, could one deny it by (O4), that Sparky running a program is not suitably an indissoluble highly-interconnected unit? Yes—because the different programs that Sparky runs are unrelated and unconnected to each other. Suppose, however, that they *were* all connected in some suitable way (say, Sparky is running “Lotus 1-2-3- ∞ ”, which treats as a single humongous unit a simultaneously running calculator, word-processor, spreadsheet, database, e-mail program, etc., etc., ad infinitum, with all data freely shared among the modules). Then, *perhaps*, Sparky could be said to think—but this is far different from arbitrary computation and far different from ordinary PCs.

8 SUMMARY.

So my pocket calculator *is* a thinking thing *if* ‘think’ is taken in the minimal sense and ‘calculate’ means “do arithmetic”. But that’s not very astounding. It is similar to the claim that thermostats have beliefs. I think that they do (in contrast, say, to Stich 1983), but it’s a very weak sense of ‘belief’.

What does the conclusion that a pocket calculator minimally thinks tell us about maximal thinking? One lesson, perhaps, is that thinking is not an all-or-nothing affair—it’s not the case that something either thinks or else it doesn’t: There are degrees of thinking, with humans at one end of the spectrum and pocket calculators nearer the other end. A goal of AI is to produce computer programs that, when executed, will be able to think at the human end of the spectrum.

How far can we travel towards the other end of the spectrum and still maintain that there is thinking going on? Suppose that Cal Jr. can only add (but not subtract, multiply, or divide); can Cal Jr. think? By Hauser’s argument, yes, if adding is (part of) calculating and calculating is (part of) thinking. Suppose MiniCal can only add 2 and 3; what then? Here, perhaps, one’s intuitions falter, but an answer consistent with Hauser’s argument is to bite the bullet and say that MiniCal can think. What about what we might call ‘MicroCal’, the “calculator”, attributed to the computer scientist Edsger Dijkstra, consisting of a piece of cardboard with ‘ $2 + 3 =$ ’ on one side and ‘5’ on the other? What about NanoCal, a piece of paper with ‘ $2 + 3 = 5$ ’ inscribed on it? What about a book of mathematical tables? What about the text of Cal’s program? What about an encyclopedia?⁵ I don’t know. I’m tempted to say that we’ve crossed a boundary, though I’d be hard-pressed to say just where. (I have an idea, though: the boundary is crossed when we move from *dynamic* systems—physical devices *executing* programs—to *static* ones—*texts* of programs, reference books, etc.)

Here’s a similar question: Are viruses alive? Humans are clearly alive. So are amoebae. Life was once thought to be something exhibited only by things endowed with a “vital spirit”; it’s now known to be (merely) very complex chemistry. Are crystals alive? (Must life be *biochemistry*?) Is “artificial life” life (cf. Meyer & Wilson 1991)? These questions are, I think, analogous to those raised by Hauser’s argument. They are worth pondering.

But my pocket calculator is *not* a thinking thing if ‘think’ is taken in the maximal sense, because *mere* calculating isn’t maximal thinking. Mere calculating can be “promoted” to maximal thinking by endowing Cal with more mental abilities. Some say that this cannot be done, or would not suffice. Perhaps; but that’s another story. And, anyway, that’s how *we* do it.

9 ACKNOWLEDGMENTS.

These comments have benefitted greatly from comments by Neal Goldberg, Bharat Jayaraman, Deepak Kumar, Stuart C. Shapiro, and the other members of the SNePS Research Group.

⁵On this example, see Lenat & Feigenbaum 1991, Smith 1991.

10 REFERENCES.

1. Crimmins, Mark (1989), "Having Ideas and Having the Concept," *Mind and Language* 4: 280–294.
2. Dretske, Fred (1985), "Machines and the Mental," *Proceedings and Addresses of the American Philosophical Association* 59: 23–33.
3. Fodor, Jerry (1983), *The Modularity of Mind* (Cambridge, MA: MIT Press).
4. Galbraith, Mary, & Rapaport, William J. (eds.) (1991), "Where Does *I* Come From?," *Technical Report No. 91-01* (Buffalo: SUNY Buffalo Center for Cognitive Science).
5. Hauser, Larry (1991), "Why Isn't My Pocket Calculator a Thinking Thing?," paper presented at the Colloquium on Philosophy of Mind, American Philosophical Association Central Division, Chicago, 27 April 1991; abstract in *Proceedings and Addresses of the American Philosophical Association* 64.6 (February 1991) 69. In L. Hauser & W. J. Rapaport (1991), "Why Isn't My Pocket Calculator a Thinking Thing? Essay, Comments, and Reply," *Technical Report 91-20* (Buffalo: SUNY Buffalo Department of Computer Science); page references are to the technical report.
6. Kenny, Anthony (1968), *Descartes: A Study of His Philosophy* (New York: Random House).
7. Lenat, Douglas B., & Feigenbaum, Edward A. (1991), "On the Thresholds of Knowledge," *Artificial Intelligence* 47: 185–250.
8. Maloney, J. Christopher (1987), "The Right Stuff," *Synthèse* 70: 349–372.
9. Marcotty, Michael, & Ledgard, Henry (1986), *Programming Language Landscape: Syntax, Semantics, and Implementation*, 2nd edition (Chicago: Science Research Associates).
10. Meyer, Jean-Arcady, & Wilson, Stewart W. (eds.) (1991), *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior* (Cambridge, MA: MIT Press).
11. Pylyshyn, Zenon (1985), *Computation and Cognition: Toward a Foundation for Cognitive Science*, 2nd edition (Cambridge, MA: MIT Press).
12. Rapaport, William J. (1988), "Syntactic Semantics: Foundations of Computational Natural-Language Understanding," in J. H. Fetzer (ed.) *Aspects of Artificial Intelligence* (Dordrecht, Holland: Kluwer Academic Publishers): 81–131.
13. Rapaport, William J. (1990a), "Computer Processes and Virtual Persons: Comments on Cole's 'Artificial Intelligence and Personal Identity'," *Technical Report 90-13* (Buffalo: SUNY Buffalo Dept. of Computer Science, May 1990).
14. Rapaport, William J. (1990b), "Cognitive Science," in A. Ralston & E. D. Reilly (eds.), *Encyclopedia of Computer Science and Engineering*, 3rd edition (New York: Van Nostrand Reinhold, forthcoming); longer version pre-printed as *Technical Report 90-12* (Buffalo: SUNY Buffalo Dept. of Computer Science, May 1990).
15. Russell, Bertrand (1950), *History of Western Philosophy* (London: Allen and Unwin).

16. Searle, John R. (1980), "Minds, Brains, and Programs," *Behavioral and Brain Sciences* 3: 417–457.
17. Smith, Brian Cantwell (1991), "The Owl and the Electric Encyclopedia," *Artificial Intelligence* 47: 251–288.
18. Stich, Stephen P. (1983), *From Folk Psychology to Cognitive Science: The Case against Belief* (Cambridge MA: MIT Press).

THE SENSE OF ‘THINKING’: Reply to Rapaport’s Comments

Larry Hauser

It will be found that the great majority, given the premiss that thought is not distinct from corporeal motion, take a much more rational line and maintain that thought is the same in the brutes as in us, since they observe all sorts of corporeal motions in them, just as in us. And they will add that “the difference, which is merely one of degree, does not imply any essential difference”; from this they will be quite justified in concluding that, although there may be a smaller degree of reason in the beasts than there is in us, the beasts possess minds which are of exactly the same type as ours. (Descartes 1642: 288–289)

1 CLARIFICATIONS AND ISSUES.

I begin with several clarifications. What I mean by ‘thinking’ is having (some) mental properties: perceptual properties (e.g., seeing, hearing, detecting), cognitive properties (e.g., knowing, believing, calculating), conative properties (e.g., wanting, needing, seeking), etc. ‘Calculate’, in the context of this argument, means doing arithmetic: adding, subtracting, multiplying, dividing, determining percentages, and extracting square roots; Cal’s apparent abilities. The major premise, ‘Calculating is thinking’, is meant to assert that whatever calculates thinks, not vice versa. My cat, Mary Jane, for instance, certainly knows where her food bowl is and often wants to go out, but M.J. can’t add, subtract, multiply, or divide.

This last remark broaches related questions which seem the central points of contention between Rapaport and me: (1) whether the argument I present shows that Cal thinks in a philosophically interesting sense, and (2) whether there is a distinction of *meaning* or of *sense* between what Rapaport calls the “minimal” and “maximal” senses of ‘thinking’.

2 THE PHILOSOPHICAL INTEREST OF CAL’S CLAIM TO BE THINKING.

With respect to the first question, I note that Searle and Dretske both explicitly deny that pocket calculators really calculate—that they really add, subtract, multiply, or divide, or can really (i.e., literally and truly) be said to have any mental abilities whatever. Both deny that Cal thinks even in the “minimal” sense of having some (even if just one) mental ability. Descartes would also deny that Cal calculates or has any mental abilities whatever (that Cal thinks, even in a “minimal sense”). Descartes, of course, would also deny that my cat M.J. has any mental abilities whatever (that she thinks, even in a “minimal sense”).

Rapaport wonders, “On what basis could one possibly deny that Cal was calculating?” (Rapaport 1991: 15). Dretske and Searle would say, “Cal lacks intrinsic intentionality” (the intentionality objection). Descartes, if he were to put it in Rapaport’s terms, might say, “Only *maximal* thinking is real thinking, so piecemeal thought-like abilities, such as Cal’s so-called ‘calculating’ and M.J.’s so-called ‘desire to go out’ are not really thinking (not real modes of thought) at all.” I share Rapaport’s amazement that anyone *would* deny, with Searle and Dretske, that calculators calculate—I find it as amazing as Descartes’s denial that cats, e.g., really see and

hear and want things.¹ I believe that opponents of AI are driven to say such things in order to save the thesis that machines don't think, and the reason they are driven to such lengths to maintain that thesis is the difficulty of making out any such distinction of *meaning* between "maximal thinking" and "minimal thinking" as Rapaport tries to make out.

3 UNAMBIGUOUS 'THINKING'.

I turn then, to Rapaport's attempt to make out such a "maximal sense" of 'thinking'—a sense he roughly identifies with thinking to a degree that suffices, or having a sufficient number and range of mental abilities, to enable you to pass Turing's Test. And the first thing to note is that if passing Turing's Test is the criterion of thinking in the same (maximal) sense as you and I, then neither cats, nor dogs, nor chimps, nor porpoises nor any other infrahuman animal thinks in the same sense that humans do. So if the question at issue is whether Cal belongs to the "club" of things with minds, along with us, and you think that this "club" includes cats and frogs (as Searle admits) and perhaps grasshoppers (Searle says he's not sure about them), then the Turing Test criterion (and likewise, perhaps, any other holistic criterion or unity criterion) is going to rule out Cal at the cost of ruling out other things (e.g., M.J.) that we may want to include. Perhaps this is why Dretske and Searle appeal to intentionality and not unity to rule out Cal, and Descartes, who does appeal to the unity of minds or universality of mental capability, takes this holistic turn, among other things, to exclude cats (whales, monkeys, clams, sponges, etc.) from the "club".

I admit that the notion of "various sorts" of mental abilities is not entirely clear. Neither is it clear which sorts and how many (of what sorts) are required for "maximal thinking"—which is why I only say (p. 6) that some holistic requirement "*may* suffice to rule out Cal" (which is why I haven't seriously undercut my own argument). Well, some things have more mental properties than others. I trust I have more mental properties than my cat M.J., for instance, but I suppose that both she and I have mental properties (can be said to "think") in the same sense. Similarly, Lake Michigan contains more water than this glass contains—you might say that Lake Michigan contains water maximally and the glass contains water minimally—yet surely both *contain water* in the same sense. More generally, the notion that there are "degrees of thinking with humans at one end of the spectrum and pocket calculators nearer the other end" (Rapaport 1991: 17) squares ill with the claim that 'thinks' means something different when I say "Cal thinks" (because he calculates, and calculating is thinking) than when I say "You think" (because you hear me speak, agree with some of my points, disagree with others, etc.); it squares ill with the claim that the difference between my thinking and Cal's thinking answers to a difference in the sense of the word 'thinking'.

If the preceding is correct, then Rapaport has not shown—if such was his intention—that my argument equivocates on the term 'thinks'. He doesn't establish that Cal doesn't think in the same *sense* (though not, of course, to the same extent—who would say that?) as you and I. "An ordinary application of Occam's razor places the onus of proof on those who wish to claim that these sentences are ambiguous. One does not multiply meanings beyond necessity." (Searle 1975: 40.)

¹Much as Searle would say that what Cal does is just (unconscious, mechanical) "as-if" calculating, Descartes maintains that lower animals just have (unconscious, mechanical) "as-if" desires, sensations, etc.

REFERENCES

1. Descartes, René, (1642) “Replies to the Sixth Set of Objections,” trans. in J. Cottingham, R. Stoothoff, D. Murdoch, *The Philosophical Writings of Descartes*, Vol. 2 (Cambridge, Eng.: Cambridge University Press, 1984): 285–300.
2. Rapaport, William J. (1991), “Because Mere Calculation Isn’t Thinking: Comments on Hauser’s ‘Why Isn’t My Pocket Calculator a Thinking Thing?’,” presented orally at the Colloquium on Philosophy of Mind, American Philosophical Association Central Division, Chicago, 27 April 1991. In L. Hauser & W. J. Rapaport (1991), “Why Isn’t My Pocket Calculator a Thinking Thing? Essay, Comments, and Reply,” *Technical Report 91-20* (Buffalo: SUNY Buffalo Department of Computer Science); page references are to the technical report.
3. Searle, John R. (1975), “Indirect Speech Acts,” in *Expression and Meaning* (Cambridge, Eng.: Cambridge University Press, 1979): 30–57.
4. Turing, Alan M. (1950), “Computing Machinery and Intelligence,” *Mind* 59: 433–460.; reprinted in A. R. Anderson (ed.), *Minds and Machines* (Englewood Cliffs, NJ: Prentice-Hall, 1964): 4–30.