communicating moves has been established. Meanwhile a third man—B—is operating a machine that has been programmed to play chess. C plays a game with either A or the machine operated by B. Will he be able to guess which one is his opponent? Turing suspects that he will find it quite difficult to tell the difference and concludes by remarking parenthetically that the experiment is one he has actually performed. He does not give the result, however, and so "Intelligent Machinery" ends with a number of questions hanging: can a machine, educated through a system of reward and punishment, be said to be able to think? Are children, when they cry or laugh, revealing some spark of soul that distinguishes them from machines, or simply following "rules of behavior" with which we as spectators empathize because we are familiar with them? Or to put it another way, does asking whether computers think require us to ask, as well, whether humans compute?

7

The Imitation Game

1.

The Manchester in which Alan Turing settled in the fall of 1948 was as noteworthy for its industrial ugliness as for its bad weather. Manchester University, just outside the city center, was equally depressing. In Newman's laboratory, the walls were covered with brown tiles in what F. C. Williams, his partner in the project, called a "late lavatorial" style. Most of the faculty lived in the suburb of Hale, where Turing rented rooms before buying his first and only house, in 1950, on Adlington Road in Wilmslow, Cheshire. These rooms likely resembled one that W. G. Sebald described in *The Emigrants*, "carpeted in a large floral pattern, wallpapered with violets, and furnished with a wardrobe, a washstand, and an iron bedstead with a candlewick bedspread."

The machine on which Turing went to work was a preliminary model intended for small-scale experiments, and thus christened (in keeping with Turing's educational program) the Baby. It had the distinction, however, of employing Williams's and Kilburn's cathode-ray tube technology, which meant that for the first time both the instructions fed into the machine and the results it spit out could be *seen*. Not that the Baby employed anything so sophisticated as a screen: instead, the numbers appeared in form of bright spots on the monitor tubes themselves. Spots, or "bits," were arranged on each tube in a 32 × 32 grid (for a total of 1,024 bits), with each bit charged to represent a 0 or 1. A metal pickup plate was set up to detect the charge, and thus "read" the bit's value. Each 32-bit line in the grid, in turn, represented either a number or an instruction; later, the lines would be lengthened to 40 bits each, with each addressable line containing either one 40-bit number or two 20-bit instructions. As Turing remarked in the programmer's handbook that he prepared for the Manchester computer, the information in the electronic store could be compared "to a number of sheets of paper exposed to the light on a table, so that any particular word or symbol becomes visible as soon as the eye focusses"—an analogy that recalls the perforated sheets employed at Bletchley in the effort to break the Enigma code.

One of the oddities of working with the Manchester computer was the programming notation with which, as Martin Campbell-Kelly puts it, Turing "saddled users of the machine. . . . Each program instruction consisted of 20 bits, which Turing wrote down as four characters using the 5-bit Post Office teleprinter code. In effect he used the teleprinter code as a base-32 number system. . . ." This in turn required Turing to invent a 32-symbol "alphabet" of number equivalencies in which most numbers were paired with letters—9 was D, for instance; 19 was W—while some were represented by symbols (@ for 2, " for 27, £ for 31) and 0 was represented by a slash (/). "Because zero was represented by the forward-stroke character," Campbell-Kelly explains, "and this was the most commonly-used character in the written form of programs and data, one early user decided this must be an unconscious reflection of the famously dismal Manchester weather as the effect was that of rain seen through a dirty window pane!" (//////////////) As if things weren't complicated enough, numbers entered into the machine had to be written backward. Using the base-32 code, the 40-digit binary sequence 10001 11011 10100 01001 10001 11001 01010 10110 (in denary notation, 17 27 5 18 17 19 10 13) would thus have to be written as Z"SLZWRF—which would, of course, first have to be reversed. This had the effect of leaving anyone who wished to use the machine—including Turing's assistants, Audrey Bates and Cicely Popplewell—rather beholden to its language teacher. Indeed, when Turing delivered a lecture on "Checking a Large Routine" at Cambridge on June 24, 1949 (the day after his thirty-seventh birthday), his failure to bother to clarify the notational system in which he was writing figures on the blackboard struck Maurice Wilkes, who was in the audience, as "bizarre in the extreme. . . . [Turing] had a very nimble brain himself and so no need to make concessions to those less well-endowed." The base-32 code was rather like the bicycle that Turing had had at Bletchley, rigged up so that no one but he could ride it.

By way of an experiment to test the efficiency of the Baby, Newman decided to put to it one of the great puzzles of pure mathematics. This involved the so-called Mersenne primes, named after the French monk Marin Mersenne (1588–1648), who in 1644 undertook an investigation into the interesting fact that certain large prime numbers take the form $2^n - 1$ where n is also prime. As Mersenne soon discovered, the rule did not hold for all prime *n*'s. (For instance, 2" – 1 isn't prime, though 11 is.) However, by the nineteenth century it had been shown that the rule did hold when *n* was equal to 2, 3, 5, 7, 13, 19, 31, 67, and 127. In 1876 Edouard Lucas (1842–1891) came up with a method by which $2^{127} - 1$ was shown to be prime, and in 1932 D. H. Lehmer (1905–1991) was able to establish that $2^{257} - 1$ was *not* prime. Subsequently, the Mersenne numbers up to $2^{521} - 1$ were found to be not prime. A number as huge as $2^{521} - 1$, Newman realized, was probably beyond the Baby's scope; his objective, however, was less to make a discovery than to assess the computer's capacities. Accordingly, he set the baby to the task of testing Mersenne primes, using Lucas's method, which required it first to divide the numbers in question into blocks of 40 digits each and then to program the necessary carrying. In the end, though it found no new primes, the Baby was able to verify both Lucas's and Lehmer's findings—no mean feat, and a good indication of its potential.*

Operating the Manchester machine wasn't easy. Among other tasks, the operator had frequently to run from the machine room to the tape room upstairs, where the engineer would, on her instructions, switch the writing current on and then off again. A great amount of physical energy had to be expended, and there was vast room for error. "As every vehicle that drove past was a potential source of spurious digits," Cicely Popplewell later recalled, "it usually took many attempts to get a tape in—each attempt needing another trip up to the tape room." Indeed, the members of the Manchester team were soon so lost in the technical complexities of actually getting the machine to do its job that when news of their

*Julia Robinson later proved that $2_{521} - 1$ was, in fact, prime.

The Manchester Computer in 1955. (© Hulton-Deutsch Collection/CORBIS)

research reached the press, they were ill-prepared to deal with the consequences. And as it happened, the 1948 publication of a book called *Cybernetics*, by the American Norbert Wiener (1894–1964), had started a chain of events that cast upon the Manchester project an unwanted spotlight.

What happened was this: Wiener, who admired Turing, made a special trip to visit him in the spring of 1947 in order to discuss the future of intelligent machines. Wiener's writings were much more sensationalistic than Turing's, in addition to which he was something of a futurist manqué, inclined to play up (for instance) the similarity between nerves and electrical circuits and to prophesy scenarios in which robots working at factories render their human counterparts redundant.

Word of Wiener's ideas and his visit soon reached the ears of Sir Geoffrey Jefferson (1886–1961), the chair of the Department of Neurosurgery at Manchester University and an early advocate of the frontal lobotomy. Jefferson was due to give the Lister Oration at Manchester on June 9, 1949, and chose as his topic "The Mind of Mechanical Man." In effect, the purpose of the speech was to expose and debunk the Manchester computer project, while hymning the innate superiority of the human soul to anything mechanical or man-made:

> Not until a machine can write a sonnet or compose a concerto because of thoughts and emotions felt, and not by the chance fall of symbols, could we agree that machine equals brain—that is, not only write it but know that it had written it. No mechanism could feel (and not merely artificially signal, an easy contrivance) pleasure at its success, grief when its valves fuse, be warmed by flattery, be made miserable by its mistakes, be charmed by sex, be angry or miserable when it cannot get what it wants.

In his report for the NPL, Turing had also addressed, in a rather tongue-in-cheek way, the claim that even if provided with a method of locomotion and sense organs, a machine would still be incapable of enjoying much of what human

beings enjoyed. For Turing, however, this was of no consequence: as he later put it, the ability to enjoy strawberries and cream was not a prerequisite for intelligence. Jefferson, on the other hand, brandished the machine's supposed lack of consciousness as evidence of its ultimate stupidity. Summarizing his oration the next day, the *Times* of London paraphrased him as saying that unless a machine could "create concepts and find for itself suitable words in which to express them . . . it would be no cleverer than a parrot"; the paper also reported that Jefferson "feared a great many airy theories would arise to tempt them against their better judgment, but he forecast that the day would never dawn when the gracious rooms of the Royal Society would be converted into garages to house the new fellows."

This was clearly meant as a slight to Newman, whose project the Royal Society had funded, and a day later the newspaper followed up with an article on Newman's "mechanical brain," noting that the "mechanical mind" had "just completed, in a matter of weeks, a problem, the nature of which is not disclosed, which was started in the seventeenth century and is only just being calculated by human beings." The machine was described as being "composed of racks of electrical apparatus consisting of a mass of untidy wires, valves, chassis, and display tubes. When in action, the cathode ray becomes a pattern of dots which shows what information is in the machine. There is a close analogy between its structure and that of the human brain." The article also included an interview with Turing, who said of the machine,

> This is only a foretaste of what is to come, and only the shadow of what is going to be. We have to have some expe-

> rience with the machine before we really know its capabilities. It may take years before we settle down to the new possibilities, but I do not see why it should not enter any of the fields normally covered by the human intellect and eventually compete on equal terms.
>
> I do not think you can even draw the line about sonnets, though the comparison is perhaps a little bit unfair because a sonnet written by a machine will be better appreciated by another machine!

There was no reason to assume, in other words, that even poetry (Jefferson had ended his oration by quoting from *Hamlet*) should be the exclusive province of the human imagination. (A relative who read the article told Mrs. Turing, "Isn't that just like Alan?") Yet what is more striking than Turing's willingness to attribute to a machine the capacity for writing and understanding verse is his suggestion that machines might speak between themselves a language no less meaningful for its exclusion of human beings. It was as if what offended Turing, even more than Jefferson's avidity to shut down avenues of exploration, was his hawking of "humanist" values for the explicit purpose of denying a whole class of beings the right to a mental existence. Likewise homosexual men, for decades, had been erased from history—and more specifically, from the history of human eros to which Jefferson alluded by mentioning "the charm of sex." In any case, Turing told the *Times*, "The university was really interested in the investigation of machines for their own sake." It was as if, by this point, he was becoming sick of the human.

As for Newman, he gave his own reply to the *Times* in the form of a letter published on June 14, in which he attempted

to summarize some of the science behind the prototype Manchester machine and also clear the air regarding "the rather mysterious description" that the newspaper had given of the problem dating back to the seventeenth century. Testing out the Mersenne primes, he explained, was exactly the sort of pure mathematical exercise at which Newman hoped his machine would excel. Indeed, the earnestness with which he attempted to make the experiment comprehensible to the *Times*'s readers provided clear evidence as to just how far apart the perspective of the Manchester laboratory was from the one that informed Jefferson's oration. Nonetheless, the letters column of the *Times* continued, for a few days, to offer evidence that perhaps Turing and Newman were underestimating the hostility that their research had the potential to provoke. England was as disinclined to accept machine nature as human nature, if Illtyd Trethowan of Downside Abbey, Bath, was to be believed; in a letter to the *Times* dated June 13, he expressed his hope that "responsible scientists will be quick to dissociate themselves" from Newman's program. "But we must all take warning from it. Even our dialectical materialists would feel necessitated to guard themselves, like Butler's Erewhonians, against the possible hostility of the machines."*

*Turing might have been thinking of Trethowan when at the end of his 1951 Manchester lecture he remarked, "There would be plenty to do in trying, say, to keep one's intelligence up to the standard set by the machines, for it seems probable that once the machine thinking method had started, it would not take long to outstrip our feeble powers. There would be no question of the machines dying, and they would be able to converse with each other to sharpen their wits. At some stage therefore we should have to expect the machines to take control, in the way that is mentioned in Samuel Butler's *Erewhon*."

As for Jefferson's remark that unless a machine could "create concepts and find for itself suitable words in which to express them . . . it would be no cleverer than a parrot," it provoked a spirited defense of the bird in the paper's editorial pages that brought an end to the brouhaha, with the author facetiously complaining that

> those who have never loved a parrot can hardly appreciate the vehemence of the emotions aroused by these thoughtless words in the breasts of those who have made of this sagacious bird a close and (as far as can be ascertained) devoted companion. . . . Parrots can make things devilish unpleasant if they take a dislike to you, and it would be a prudent as well as a courteous gesture if Professor Jefferson withdrew an observation which has ruffled so many and so well-loved feathers.

So far as the *Times* was concerned, the call for an apology to the parrot (but not to the scientists) brought the matter to a close. But Turing did not forget what Jefferson had said. If anything, the exchange in the newspaper's pages only strengthened his interest in machine intelligence. He would soon strike back—and even the parrot would make another appearance.

2.

"Computing Machinery and Intelligence," Alan Turing's most famous and in many ways most perverse paper, appeared in *Mind* in October 1950. Whereas in the NPL report he started with likely objections, here he saved the list of potential objections to computer intelligence for later, and began instead

with a clear statement of his intent. "I propose to consider the question, 'Can machines think?' This should begin with definitions of the meaning of the terms 'machine' and 'think.' " But if these meanings "are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, 'Can machines think?' is to be sought in a statistical survey such as a Gallup poll." Such an idea, in Turing's view, was "absurd."

Instead of offering definitions, Turing recast his question by proposing what he called the imitation game. It would later become known as the Turing test, much as the *a*-machine of "Computable Numbers" has come to be called a Turing machine. The game, as he explains it,

> is played with three people, a man (A), a woman (B), and an interrogator (C) who may be of either sex. The interrogator stays in a room apart from the other two. The object of the game for the interrogator is to determine which of the other two is the man and which is the woman. He knows them by labels X and Y, and at the end of the game he says either "X is A and Y is B" or "X is B and Y is A." The interrogator is allowed to put questions to A and B thus:
>
> C: Will X please tell me the length of his or her hair?
>
> Now suppose X is actually A, then A must answer. It is A's object in the game to try and cause C to make the wrong identification. His answer might therefore be
>
> "My hair is shingled, and the longest strands are about nine inches long."
>
> In order that tones of voice may not help the interrogator the answers should be written, or better still, type-

written. The ideal arrangement is to have a teleprinter communicating between the two rooms. Alternatively the question and answers can be repeated by an intermediary. The object of the game for the third player (B) is to help the interrogator. The best strategy for her is probably to give truthful answers. She can add such things as "I am the woman, don't listen to him!" to her answers, but it will avail nothing as the man can make similar remarks.

We now ask the question, "What will happen when a machine takes the part of A in this game?" Will the interrogator decide wrongly as often when the game is played like this as he does when the game is played between a man and a woman? These questions replace our original, "Can machines think?"

Turing's proof, in "Computable Numbers," that the *Entscheidungsproblem* was insoluble relied on the ingenious substitution of a complicated question—can a machine decide whether a statement is provable?—with a simpler one: does a certain machine ever print a 0? Along the same lines, in "Computing Machinery and Intelligence" he argued that the complicated question "Can machines think?" could be substituted with the simpler question "Can a machine win the imitation game?" The two, in Turing's view, were identical, because behavior, as he saw it, *was* identity. And yet to apply such a mathematically precise notion of identity to the murky matter of what "human" meant was to invite all sorts of objections—and problems.

For example, the ambiguity of Turing's query "What will happen when the machine takes the part of A?" has occasioned much debate. Does Turing mean to say that instead of being played between a man and a woman, the game should be

played between a man and a machine? The rest of the paper would seem to bear out this interpretation. Yet a literal reading of the paragraph suggests a different meaning: that the game should now be played between a man *and a computer pretending to be a man pretending to be a woman*. Hodges shows little patience for this reading, going so far as to argue that "Turing's gender-guessing analogy detracts from his own argument. . . ." After all, as he points out, the section that follows the troublesome paragraph is entirely concerned with the ways in which a machine might trick an interrogator into believing that he (or she) was talking to a human being—male *or* female:

> The new problem has the advantage of drawing a fairly sharp line between the physical and the intellectual capacities of a man. No engineer or chemist claims to be able to produce a material which is indistinguishable from the human skin.* It is possible that at some time this might be done, but even supposing this invention available we should feel there was little point in trying to make a "thinking machine" more human by dressing it up in such artificial flesh. The form in which we have set the problem reflects this fact in the condition which prevents the interrogator from seeing or touching the other competitors, or hearing their voices. Some other advantages of the proposed criterion may be shown up by specimen questions and answers. Thus:
>
> Q: Please write me a sonnet on the subject of the Forth Bridge.

*For an interesting analysis of skin imagery—of which there is a lot—in Turing's paper, see Jean Lassègue, "What Kind of Turing Test Did Turing Have in Mind?" *http://tekhnema.free.fr/3Lasseguearticle.htm.*

A: Count me out on this one. I never could write poetry.
Q: Add 34957 to 76764.
A: (Pause about 30 seconds and then give as answer) 105621.
Q: Do you play chess?
A: Yes.
Q: I have K at my K1, and no other piece. You have only K at K6 and R at R1. What do you play?
A: (After a pause of 15 seconds) R-R8 mate.

Hodges is correct to observe that gender plays no role in the answers given here (including the incorrect addition). And yet to ignore the subtext that Turing's ambiguity exposes is also to ignore the palpable tone of sexual anxiety that runs all through the paper. For instance, just a few paragraphs after the dialogue quoted above, Turing writes, "It might be urged that when playing the 'imitation game' the best strategy for the machine may possibly be something other than imitation of the behaviour of a man. . . . In any case there is no intention to investigate here the theory of the game, and it will be assumed that the best strategy is to try to provide answers that would naturally be given by a man."* Trying to provide "answers that would naturally be given by a man" would, of course, also be the best strategy for a homosexual to adopt when trying to persuade an interrogator that he is straight; in this alternative version of the imitation game, he would talk about cricket and

*Turing used similar language during a 1952 BBC roundtable discussion, in which, as an example of the sort of question to use in the imitation game, he proposed the following: "I put it to you that you are only pretending to be a man." In such a case, "the machine would be permitted all sorts of tricks so as to appear more manlike. . . ."

describe the woman he would like to marry. And though the parallel may be accidental—"a man," after all, could as easily mean "a human being" as "a male human being"—Turing's use of the word "naturally" suggests a more heightened awareness of the idea of the "natural" than the situation calls for. Not surprisingly, arguments concerning the naturalness or unnaturalness of homosexuality ran through both antihomosexual diatribes and apologies for homosexuality written in the period, with Oscar Wilde's championing of the artificial often brandished as an ironic defense of "unnatural" love.*

Turing's preoccupation with gender recurs several more times during the course of the paper. In section 3, a discussion of exactly what defines a "machine" concludes with this rather bizarre proviso that

> we wish to exclude from the machines men born in the usual manner.
>
> It is difficult to frame the discussion so as to satisfy [this condition]. One might for instance insist that the team of engineers should be all of one sex, but this would not really be satisfactory, for it is probably possible to rear a complete individual from a single cell of the skin (say) of a man. To do so would be a feat of biological technique deserving of the very highest praise, but we would not be inclined to regard it as a case of "constructing a thinking machine."

Is the point here that the team of engineers—all of "one sex"—might be able to join together and in a sort of orgy of

*In *Maurice* the hero asks Alec, "Scudder, why do you think it's 'natural' to care both for men and women? You wrote so in your letter. It isn't natural for me. I have really got to think that 'natural' only means oneself."

cloning create a human child? The fantasy is peculiar, using science as a framework for imagining a means by which men without women could generate progeny. Of course, Turing too longed to produce a child of his own—a computer child. It is therefore not surprising that in the paper he soon returns to the metaphor of child rearing and education, employing a "domestic analogy" to describe the ways in which a machine might be taught to obey "not fresh instructions on each repetition, but the same ones over and over again":

> Suppose Mother wants Tommy to call at the cobbler's every morning on his way to school to see if her shoes are done, she can ask him afresh every morning. Alternatively she can stick up a notice once and for all in the hall which he will see when he leaves for school and which tells him to call for the shoes, and also to destroy the notice when he comes back if he has the shoes with him.

"Tommy" is here the computer, the offspring of a group of engineers who have eschewed cloning in favor of other styles of cooperation—perhaps the sort of cooperation in scientific experiment that Turing so cherished in his friendship with Christopher Morcom. More importantly, Tommy is a digital computer, and in Turing's estimation only a digital computer—a universal machine—has a shot at ever winning the imitation game.

> I believe that in about fifty years' time* it will be possible to programme computers, with a storage capacity of about

*By 1952, when he was interviewed on the BBC, the estimate had gone up to at least a hundred years.

> 10^9, to make them play the imitation game so well that an average interrogator will not have more than 70 per cent chance of making the right identification after five minutes of questioning.

The Manchester "Baby" is clearly growing up.

3.

By this point, then, a subtle but distinct strain of anxiety concerning gender, sexual imitation, and even homosexual procreation has come to assert itself within Turing's "official" argument about machine intelligence. But where does it come from? The answer can be traced back to Sir Geoffrey Jefferson's Lister Oration, the slightly masculinist tone of which Turing ridicules in the paper, even as he rebuts Jefferson's "humanist" stance. This is especially evident near the middle of "Computing Machinery and Intelligence," where Turing takes up once again the strategy of listing—and then refuting—the objections that might be raised to the possibility of a thinking machine. Although Professor Jefferson does not appear by name until the fourth objection—"The Argument from Consciousness"—his spirit is invoked, and mocked, from the very start.

For instance, in his refutation of the first objection—"the theological objection" that "God has given an immortal soul to every man and woman, but not to any other animal or to machines"—Turing questions the implicit superiority of mankind that provided the basis for Jefferson's diatribe, noting, "I should find the argument more convincing if animals were classed with men, for there is a greater difference, to my mind, between the typical animate and the inanimate than there is between man and the other animals." Likewise, how are Christians to contend with "the Moslem view that women have no souls"? By invoking the rights not just of women but of animals, Turing allies himself (and his computer) with all the other populations that have suffered at the hand of religions that take the superiority of man (in one case) and mankind (in the other) for granted. Against this he posits his own rather odd theology, which, needless to say, blesses machines, by equating their construction with procreation: "In attempting to construct such machines we should not be irreverently usurping His power of creating souls, any more than we are in the procreation of children: rather we are, in either case, instruments of His will providing mansions for the souls that He creates."

The assumption of mankind's innate superiority is challenged even more boldly in Turing's retort to the second objection, which he calls the "Heads in the Sand" objection and sums up as follows: "The consequences of machines thinking would be too dreadful. Let us hope and believe that they cannot do so." This, of course, was the very posture to which some of Norbert Wiener's writings inadvertently appealed, and in responding to it, Turing also responds to Jefferson, noting that the feeling that mankind is "*necessarily* superior" to the rest of creation "is likely to be quite strong in intellectual people, since they value the power of thinking more highly than others, and are more inclined to base their belief in the superiority of Man on this power." With his allusions to Shakespeare, Jefferson is exemplary of these "intellectual people" whose tendency to exalt their own species Turing shows so little patience for. It is a point he returns to in his

answer to the third objection, the "mathematical objection," which is essentially the argument (paraphrased in the NPL report) that his own resolution of the *Entscheidungsproblem*, in conjunction with Gödel's findings, proves "that there are certain things that . . . a machine cannot do." Turing was obviously made uncomfortable by the possibility that his solution to the *Entscheidungsproblem* might be employed in an attack on the machine that the *Entscheidungsproblem* propelled him to create. In responding to it here, however, he focuses squarely on the psychology of what might be called the natural "superiority complex" of human beings (especially intellectuals), observing shrewdly that when a machine gives a wrong answer to

> the appropriate critical question . . . this gives us a certain feeling of superiority. Is this feeling illusory? It is no doubt quite genuine, but I do not think too much importance should be attached to it. We too often give wrong answers to questions ourselves to be justified in being very pleased at such evidence of fallibility on the part of machines. Further, our superiority can only be felt on such an occasion in relation to the one machine over which we have scored our petty triumph. There would be no question of triumphing simultaneously over *all* machines.

Here Turing seems to be amusing himself, in a rather quiet way, by alluding to Mr. Illtyd Trethowan's anxiety about "the possible hostility of the machines," over *all* of whom we can never hope to triumph. More importantly, this rebuttal gives him the chance to repeat one of his key points—that fallibility is a key ingredient in intelligence.

It is in his refutation of objection 4—"the Argument from Consciousness"—that Turing takes direct aim at Jefferson, whom he begins by quoting and at whom he hurls one of his most memorable and witty retorts:

> This argument appears to be a denial of the validity of our test. According to the most extreme form of this view the only way by which one could be sure that a machine thinks is to *be* the machine and to feel oneself thinking. . . . Likewise according to this view the only way to know that a *man* thinks is to be that particular man. It is in fact the solipsist point of view. It may be the most logical view to hold but it makes communication of ideas difficult. A is liable to believe "A thinks but B does not" whilst B believes "B thinks but A does not." Instead of arguing continually over this point it is usual to have the polite convention that everyone thinks.

Rather cleverly, Turing writes that he is "sure that Professor Jefferson does not wish to adopt the extreme and solipsist point of view." He then compares his own imitation game with a game called *viva voce*, the purpose of which is "to discover whether some one really understands something or has 'learned it parrot fashion.'" Notably, the exemplary *viva voce* that Turing cites is replete with literary references, with the questioner first asking his subject about Shakespeare, then veering into Dickens. The point is that the imitation game *also* determines whether someone has learned something "parrot fashion"; it differs from *viva voce* only in that the person being tested is a machine. Nor is it a coincidence that literature plays such a prominent role in this *viva voce*, the orchestrator of which is presumably a self-proclaimed intellectual such as Jefferson. And surely any self-respecting intellectual would rather abandon the argument from consciousness "than be forced into the solipsist position."

Having got rid of Jefferson—at least in name—Turing next addresses a whole class of objections that he calls "Arguments from Various Disabilities," and which he defines as taking the form "I grant you that you can make machines do all the things you have mentioned but you will never be able to make one to do X." He then offers a rather tongue-in-cheek "selection":

> Be kind, resourceful, beautiful, friendly; have initiative, have a sense of humour, tell right from wrong, make mistakes; fall in love, enjoy strawberries and cream; make some one fall in love with it, learn from experience; use words properly, be the subject of its own thought; have as much diversity of behaviour as a man, do something really new.

As Turing notes, "no support is usually offered for these statements," most of which are

> founded on the principle of scientific induction. . . . The works and customs of mankind do not seem to be very suitable material to which to apply scientific induction. A very large part of space-time must be investigated, if reliable results are to be obtained. Otherwise we may (as most English children do) decide that everybody speaks English, and that it is silly to learn French.

Turing's repudiation of scientific induction, however, is more than just a dig at the insularity and closed-mindedness of England. His purpose is actually much larger: to call attention to the infinite regress into which we are likely to fall if we attempt to use disabilities (such as, say, the inability, on the part of a man, to feel attraction to a woman) as determining factors in defining intelligence. Nor is the question of homosexuality far from Turing's mind, as the refinement that he offers in the next paragraph attests:

> There are, however, special remarks to be made about many of the disabilities that have been mentioned. The inability to enjoy strawberries and cream may have struck the reader as frivolous. Possibly a machine might be made to enjoy this delicious dish, but any attempt to make one do so would be idiotic. What is important about this disability is that it contributes to some of the other disabilities, e.g. to the difficulty of the same kind of friendliness occurring between man and machine as between white man and white man, or between black man and black man.

To the brew of gender and sexuality, then, race is added, as "strawberries and cream" (earlier bookended between the ability to fall in love and the ability to make someone fall in love) becomes a code word for tastes that Turing prefers not to name. In many ways the passage recalls the rather campy bathhouse scene in the 1960 film *Spartacus,* in which a dialogue about other "delicious dishes" encodes a subtle erotic bargaining between Crassus (Laurence Olivier) and his slave Antoninus (Tony Curtis).

> *Crassus:* Do you eat oysters?
> *Antoninus:* When I have them, master.
> *Crassus:* Do you eat snails?

> *Antoninus:* No, master.
> *Crassus:* Do you consider the eating of oysters to be moral, and the eating of snails to be immoral?
> *Antoninus:* No, master.
> *Crassus:* Of course not. It's all a matter of taste.
> *Antoninus:* Yes, master.
> *Crassus:* And taste is not the same as appetite and therefore not a question of morals, is it?
> *Antoninus:* It could be argued so, master.
> *Crassus:* Um, that'll do. My robe, Antoninus. Ah, my taste . . . includes both oysters and snails.

In this exchange Crassus, too, is engaging in a kind of imitation game, the purpose of which is to assess whether it would or would not be a good idea to offer Antoninus (who prefers oysters) some of his snails. Antoninus, at the same time, recognizes the advantage, at least on occasion, of giving the "wrong" answer ("No, master")—just as a machine would have to if it were to have a chance of winning the game:

> The claim that "machines cannot make mistakes" seems a curious one. . . . I think this criticism can be explained in terms of the imitation game. It is claimed that the interrogator could distinguish the machine from the man simply by setting them a number of problems in arithmetic. The machine would be unmasked because of its deadly accuracy. The reply to this is simple. The machine (programmed for playing the game) would not attempt to give the *right* answers to the arithmetic problems. It would deliberately introduce mistakes in a manner calculated to confuse the interrogator.

"Errors of functioning," then, must be kept distinct from "errors of conclusion." Nor should it be assumed that machines are not capable of deception. On the contrary, the criticism "that a machine cannot have much diversity of behaviour is just a way of saying that it cannot have much storage capacity."

Turing wraps up his catalog of possible objections to the thinking machine with four rather curious examples. The first, which he calls "Lady Lovelace's Objection" (in reference to Byron's daughter and Babbage's muse), is that computers are incapable of "originating" anything. Instead (and here Turing quotes Lady Lovelace), "a computer can do *whatever we know how to order it* to perform." But, as Turing points out, in actual practice, machines surprise human beings all the time. Turing then rebuts the "Argument from Continuity in the Nervous System"—although it is true that a discrete-state machine cannot mimic the behavior of the nervous system, "if we adhere to the conditions of the imitation game, the interrogator will not be able to take any advantage of this difference"—and assesses the "Argument from Informality of Behavior": "If each man had a definite set of rules of conduct by which he regulated life he would be no better than a machine. But there are no such rules, so men cannot be machines." This objection Turing answers, first, by distinguishing "rules of conduct" from the "laws of behavior" by which machines are presumably regulated, then by pointing out that "we cannot so easily convince ourselves of the absence of complete laws of behaviour as of complete rules of conduct." By way of example, he describes another experiment:

> I have set up on the Manchester computer a small programme using only 1000 units of storage, whereby the machine supplied with one sixteen figure number replies with another within two seconds. I would defy anyone to learn from these replies sufficient about the program to be able to predict any replies to untried values.

The last—and most peculiar—objection that Turing takes on is the argument "from extra-sensory perception," which he prefaces with a surprisingly credulous description of telepathy, clairvoyance, precognition, and psychokinesis. Of these he remarks, "Unfortunately the statistical evidence, at least for telepathy, is overwhelming. It is very difficult to rearrange one's ideas so as to fit these new facts in." Without giving a source for this "overwhelming" evidence, Turing goes on to give the "strong" argument from ESP against a machine's winning the imitation game:

> Let us play the imitation game, using as witnesses a man who is good as a telepathic receiver, and a digital computer. The interrogator can ask such questions as "What suit does the card in my right hand belong to?" the man by telepathy or clairvoyance gives the right answer 130 times out of 400 cards. The machine can only guess at random, and perhaps gets 104 right, so the interrogator makes the right identification.

For Turing, the scenario as described opens up the "interesting possibility" of equipping the digital computer in question with a random number generator.

> Then it will be natural to use this to decide what answer to give. But then the random number generator will be subject to the psycho-kinetic powers of the interrogator. Perhaps this psycho-kinesis might cause the machine to guess right more often than would be expected on a probability calculation, so that the interrogator might still be unable to make the right identification. On the other hand, he might be able to guess right without any questioning, by clairvoyance. With E. S. P. anything may happen.

Rather than offering a refutation of this argument, Turing says only that perhaps the best solution would be to put the competitors into a "telepathy-proof room"—whatever that means. One wonders what the editors of that august scientific publication *Mind* made of this bizarre appeal to a pseudoscience as baseless, if not as pernicious, as the one on the altar of which Turing would soon be laid out, as a kind of experiment. For how could they know that years before Turing had loved a boy named Christopher Morcom, with whose spirit he had been determined to remain connected even after death?*

"Computing Machinery and Intelligence" concludes with a meditation on teaching and learning that reiterates much of the technique prescribed in "Intelligent Machinery." Here, however, Turing adds the proviso that his system of punishments and rewards does not "presuppose any feelings on the part of the machine." Moving a bit away from the rigorously behaviorist ethos that animated "Intelligent Machinery," he

*See Lassègue, "What Kind of Turing Test Did Turing Have in Mind?," for an interesting discussion of the role Christopher Morcom might have played—even subliminally—in the paper.

also reminds his readers that "the use of punishments and rewards can at best be a part of the teaching process. . . . By the time a child has learnt to repeat 'Casabianca' he would probably feel very sore indeed, if the text could only be discovered by a 'Twenty Questions' technique, every 'NO' taking the form of a blow." Less emotional techniques need to be employed as well, especially when the objective is to teach the machine to obey orders in a symbolic language.

Probably the biggest shift from "Intelligent Machinery," however, is that here Turing elects to anthropomorphize his child-machine to a much greater degree than in the earlier paper, putting more emphasis on its childishness than on its machinishness. For example, near the end of the paper, he asks, "Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child's? . . . Presumably the child-brain is something like a note-book as one buys it from the stationer's. Rather little mechanism, and lots of blank sheets." But because this notebook mind is contained within a machine body, a slightly different teaching process has to be applied to it than would be to the "normal" child:

> It will not, for instance, be provided with legs, so that it could not be asked to go out and fill the coal scuttle. Possibly it might not have eyes. But however well these deficiencies might be overcome by clever engineering, one could not send the creature to school without the other children making excessive fun of it. It must be given some tuition. We need not be too concerned about the legs, eyes, etc. The example of Miss *Helen Keller* shows that education can take place provided that communication in both direc-

tions between teacher and pupil can take place by some means or other.

One thinks of Turing as a boy, "watching the daisies grow." Does he feel some sense of identification with Helen Keller, provided (as Turing was not) with an education to suit her particular disabilities? Certainly

> the imperatives that can be obeyed by a machine that has no limbs are bound to be of a rather intellectual character. . . . For at each stage when one is using a logical system, there is a very large number of alternative steps, any of which one is permitted to apply, so far as obedience to the rules of the logical system is concerned. These choices make the difference between a brilliant and a footling reasoner, not the difference between a sound and a fallacious one.

And the ability to reason is, finally, the ultimate evidence of intelligence. If it is to be attained, however, flexibility is essential, even if "the rules which get changed in the learning process are of a rather less pretentious kind, claiming only an ephemeral validity. The reader may draw a parallel with the Constitution of the United States."

In the end, Turing believes, the goal should be to do exactly what alarms Jefferson: to construct machines that "will eventually compete with men in all purely intellectual fields." Perhaps the best way to start would be to teach the machine some "very abstract activity," such as how to play chess; or perhaps it would make more sense to provide it with "the best sense organs that money can buy, and then teach it to understand and speak English." In either case, the final note that

Turing sounds in "Computing Machinery and Intelligence" combines triumph with a certain detached self-assurance. For Turing, thinking machines are inevitable, whether we like them or not. It is as if his faith in future tolerance had once again bolstered him against the very real threat of present injustice.

4.

The years Turing spent working with the Manchester computer were marked by an increasing isolation from other people, as he became less and less interested in the computer itself and more and more involved in the experiments he was using it for. Not that he only did experiments: he also wrote a programmer's handbook in which he urged potential users of the Manchester machine to employ an almost literary sensibility in designing programs. Most of his time, though, he devoted to the application of the machine to such pure mathematical problems as constructing a new proof for the word problem for semigroups, and to working with permutation theory, which had played an important role in his code breaking at Bletchley. His colleague Christopher Strachey also taught the machine to sing "God Save the King."

Probably the experiment that meant the most to Turing, however, was the one with which he had the least success. For years he had remained fascinated by the Riemann hypothesis, which for some reason he had convinced himself had to be false. True, the machine he had tried to build with Donald MacPhail at Cambridge had ended up on the scrap heap. Yet he had never forgotten his ambition of beating Titchmarsh's record for the calculation of zeros, and still hoped he might one day be able to find a zero *off* the critical line. Toward that end, in 1943 he had published a paper entitled "A Method for the Calculation of the Zeta-function" in the *Proceedings of the Mathematical Society*. Titchmarsh, using hand methods, had shown that all the zeros up to t = 1,468 were on the critical line. Now Turing put his own method to the test. In 1953 he designed a program by means of which the Manchester computer could calculate zeta zeros using its complex base-32 code, and by means of that program he proved the validity of the Riemann hypothesis as far as t = 1,540—72 more zeros than Titchmarsh had found—before the machine broke down.

It was, as Turing ruefully noted, "a negligible advance."

8

Pryce's Buoy

1.

In the spring of 1951 Alan Turing was elected a member of the Royal Society. Among the congratulatory notes he received was one from his old antagonist Sir Geoffrey Jefferson, who wrote, "I am so glad; and I sincerely trust that all your valves are glowing with satisfaction, and signalling messages that seem to you to mean pleasure and pride! (but don't be deceived!)."

As it happened, Turing and Jefferson were destined to tangle once more. The occasion was a roundtable discussion of machine intelligence broadcast on the BBC Third Programme on January 14, 1952, in which the other participants were Max Newman and Turing's old Cambridge friend Richard Braithwaite, one of the two mathematicians who had long ago asked for offprints of "Computable Numbers." Braithwaite acted as moderator, and while the conversation did little to advance the cause of the thinking machine, at the very least it gave the speakers a chance to refine and clarify some of their positions. As always, Jefferson insisted that it was the "high