The Tide of Memory, Turning

Two dramatic presentations at a recent Cold Spring Harbor meeting have shifted the balance in a long-running debate about the precise cellular basis of memory

WHAT DO A RAT IN A MAZE, Pavlov's dogs, and a student studying for final exams have in common? The answer is that they are all storing information away in long-term memory. And that means that something, somewhere, is changing in the circuitry of their brains. Somehow the neural connections corresponding to the new knowledge are being strengthened. No one knows for sure how that strengthening happens, but many neuroscientists are placing their bets on a process called long-term potentiation, LTP for short. But just how LTP works is an open question. In fact, at a recent meeting at Cold Spring Harbor, two highly respected neuroscientists presented startling findings, based on new techniques, that threaten to turn present views of the phenomenon on their head.

For more than a decade, most researchers have been convinced that the alterations underlying long-term memory take place at the synapses, the points where signals pass from one nerve cell to the next. But that's a very general picture. Inquisitive researchers have been trying to get more specific, by asking whether the changes take place in the cell that sends the signal, or in the cell across the synapse that receives it. The search for the precise site has been the subject of some of the most intense research and debate in neuroscience in the 1980s.

In the last couple of years, an answer seemed to be emerging, with the odds favoring the postsynaptic, or receiving, cell as the site of the change. But then came Richard Tsien of Stanford University Medical Center and Charles Stevens of the Salk Institute in San Diego, who stunned a Cold Spring Harbor audience earlier this month when they announced that their teams had independently arrived at the same conclusion: a big part of the change that occurs in LTP must be presynaptic.

Both research groups employed a technique called quantal analysis, which is an effort to analyze in detail the physical signal that passes across the synapse from one neuron to the next. The signal is carried by molecules called neurotransmitters. When an electrical impulse travels along a nerve cell and arrives at a synapse, it causes the cell to release "packets" of neurotransmitter. Those packets cross the synapse and bind to receptors on the postsynaptic cell.

The binding of transmitter to the postsynaptic target cell opens channels in that cell's membrane, allowing ions to flow into the cytoplasm. If enough ions cross the membrane, an electrical signal will be fired and travel through the postsynaptic cell. The average amount of ion flow at the synapse determines that synapse's strength: the ease with which it can cause the postsynaptic cell to fire. It is by increasing such synaptic strength—at specific synapses—that memory is thought to be stored.

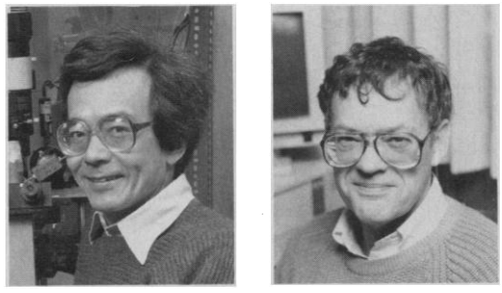

Unlocking memories. *Richard Tsien of Stanford and Charles Stevens of the Salk Institute.*

Assuming that theory is sound, the question remains of how a synapse is strengthened. There are two general possibilities. The presynaptic cell could release more packets of neurotransmitter. Alternatively, the postsynaptic cell might become more sensitive, admitting a larger flow of ions in response to the same amount of neurotransmitter.

To find out which of these possibilities is correct, Tsien and Stevens used quantal analysis to analyze the number of packets of transmitter pumped out by the sending cell. If that number were to go up after the synapse was strengthened, it would seem that (at least) a significant part of the change must occur presynaptically. And although that idea runs counter to the prevailing wisdom, that's what Tsien and Stevens found.

"I would have put money on it being the other way," says surprised University of California, San Francisco, neurobiologist Michael Stryker. With the growing evidence in favor of a postsynaptic mechanism, Stryker says, "the quantal analysis was going to be the last brick in the wall. [But] it didn't come out the way I expected."

Still, though many fascinated listeners may have thought they were seeing the issue put to rest, researchers on both sides of the debate say there are too many loose ends and unexplained observations for anyone to collect on their bets just yet.

One of the reasons feelings run so deep over LTP is that neuroscientists have been musing over the phenomenon and its possible role in memory for nearly two decades. It was in 1973 that the strengthening of synapses known as LTP was first described by Tim Bliss and his co-workers at the National Institute for Medical Research in London. Although LTP is seen elsewhere in the brain, it has been studied largely in the hippocampus, a brain area known to be involved in some types of learning and memory.

The notion that LTP plays a significant role in memory got a boost a decade later with the realization that the process required the activity of a newly discovered receptor for glutamate, a neurotransmitter used by hippocampal cells. Known as the NMDA receptor, because it can be activated by the glutamate analog NMDA, the receptor possessed special characteristics that not only offered a means of understanding associative memory—the kind Pavlov studied—but also buttressed the notion that LTP takes place on the postsynaptic, or target, side of the synapse.

Like other glutamate receptors, the NMDA receptor is activated by glutamate released by the presynaptic, or sending, cell. Unlike its counterparts, however, the NMDA receptor opens its ion channels only if the postsynaptic cell is already electrically activated—say, by a signal coming in at the same time from another nerve fiber. When such converging signals are present, NMDA receptors allow calcium ions to flow into the target cell. That influx seems to cause biochemical changes that lead to LTP.

Such a mechanism, in which two convergent signals strengthen a synapse, proved perfect for explaining associative learning. In Pavlov's experiments, two stimuli—a bell

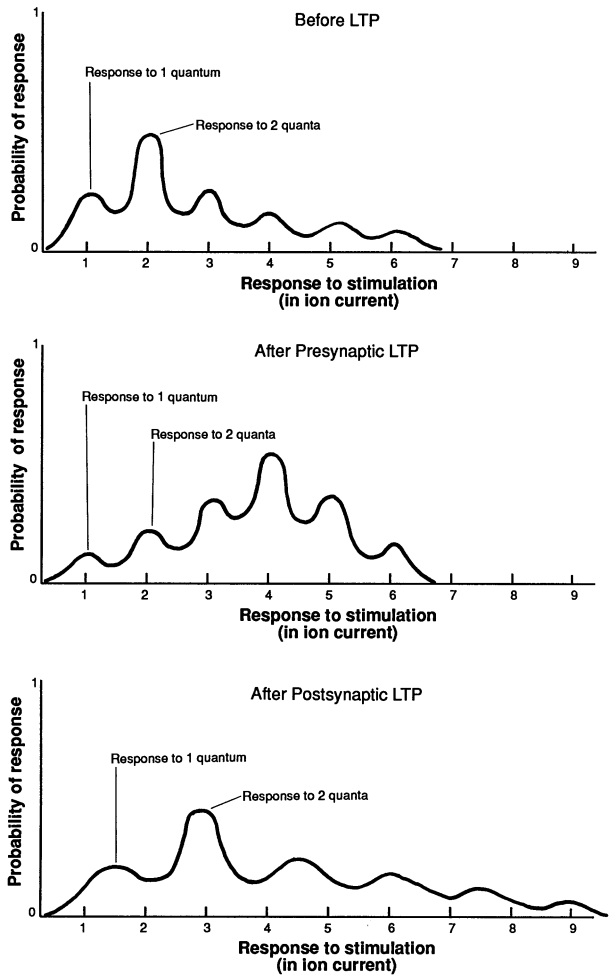

Quantum curvature. *These three idealized curves show how a technique called quantal analysis can be used to determine whether long-term potentiation (LTP) is due to presynaptic or postsynaptic modification. The y-axis of each curve shows the frequency with which a certain number of quanta (packets of neurotransmitter) are released from a presynaptic neuron. The x-axis shows the response to that number of quanta (as measured by ion flow in the postsynaptic cell). The top curve represents the control situation before LTP conditioning. The lower curves show two possibilities after LTP. If LTP were due to the release of more quanta, on average, from the presynaptic cell, the highest point of the curve would shift to the right, but the maximum response would be no higher than before (middle). If, on the other hand, LTP were due to increased sensitivity in the postsynaptic cell, the maximal response would increase (from 6 to 9), but the highest point would remain the same relative to the rest of the curve.*

and food—are initially presented together. When it sees the food, the dog salivates. But over time the dogs could be conditioned to salivate at the sound of a bell alone, demonstrating that "food" signals and "bell" signals must converge in the brain—specifically at the synapses controlling salivation. And over time, those synapses can be strengthened enough that the bell alone causes the dogs to salivate. Discovery of the NMDA receptor and its special characteristics not only offered a nifty model for this kind of learning, but also provided an argument for why the changes seen in LTP should be postsynaptic. Since LTP begins with activation of the NMDA receptors on the postsynaptic cell, it would be simplest if the synapse-strengthening changes took place postsynaptically as well. Otherwise one would have to conjure up some sort of rapid return messenger that would communicate the change back to the presynaptic cell—and this seemed too convoluted to be likely.

In spite of this buildup of evidence and presumption on the postsynaptic side, the presynaptic view was not without some adherents. Their best evidence came from Tim Bliss and his colleagues, who found that after LTP had taken place there was an increase in the amount of neurotransmitter in the synapse. That increase was taken as a sign that the presynaptic cell had begun releasing more transmitter. But other interpretations of that data were possible. And, even on the postsynaptic side, the case was far from airtight. In fact, the evidence wasn't really compelling one way or the other.

Then, last year, the postsynaptic model got strong support from independent experiments in the labs of Gary Lynch at UC Irvine and Roger Nicoll at UCSF. Those groups showed that an increase in release of glutamate by the presynaptic, or sending, cell boosted the response of both NMDA and non-NMDA glutamate receptors on the postsynaptic, or recipient, cell. In LTP, however, only the non-NMDA receptors showed such an increase. That difference, both Lynch and Nicoll argue, suggests that what is going on in LTP is quite distinct from an increase in presynaptic transmitter release—and is, in fact, much more likely to be an adjustment of receptor sensitivity in the recipient cell.

And that's roughly where things stood until Stevens and Tsien took the stage at Cold Spring Harbor. Their experiments were made possible by whole-cell recording techniques that reduce background noise to the point where very small electrical signals can be reliably detected, methods pioneered only in the last year by several groups.

Aside from that technical advance, the logic of the approach the two teams used is not new, Stevens says. Indeed, it builds on Nobel prizewinning work done more than three decades ago by Bernard Katz of University College, London, on the neuromuscular junction: the point where nerve cells release acetylcholine to cause muscle contraction. Katz identified the ion flow into a muscle cell in response to a single quantum of acetylcholine release. He also showed that the ion currents flowing into the muscle after stimulation of the nerve were the sum of the responses to individual quanta.

Katz's analysis—and those of Stevens and Tsien—depends on the probabilistic nature of a synapse. A given neuron may make thousands of synapses with its various targets. At each synapse there are a handful of sites at which packets of transmitter may be released. But no site will release transmitter every time the neuron is activated; rather, there is some probability, between 0 and 1, that a packet will be released from a particular site.

Because of this probabilistic nature, the firing of the synapse can be compared to tossing a handful of coins. Occasionally the result will be all tails: no sites, release transmitter and there is no signal in the postsynaptic cell. The result may also be all heads: all sites, release transmitter yielding the maximum signal. At other times, the response will be intermediate, with some but not all sites releasing quanta and an intermediate signal being generated.

Experimentally, this phenomenon is measured by recording the electric signal in the postsynaptic cell. If an experimenter fires the presynaptic neuron 100 times and plots the frequency of occurrence of electric signals of different strengths, the result (after mathematical modification) will be a bumpy curve, in which each bump corresponds to the average response to the release of a given number of quanta. The curve's peak will correspond to the number of quanta most frequently released into the synapse.

If the experimenter then induces LTP by stimulating the postsynaptic cell in a way that activates NMDA receptors and strengthens the synapse, subsequent stimulation of the presynaptic cell will yield, on average, a larger response in the target. But what is it that is getting bigger—the average number of quanta released, or the size of the postsynaptic response to each quantum? In other words, to return to the central ques-

tion of the debate: do the changes of LTP occur in the sending or the receiving cell?

Quantal analysis can provide an answer through examination of what happens to the bumpy curve representing the electrical signals seen in the target cell [see box on opposite page]. If the average number of quanta (that is, the amount of transmitter released) increases, the peak will shift to the right but the largest response remains unchanged. If, on the other hand, the postsynaptic cell becomes more sensitive, the curve will spread out uniformly: the bumps will get farther apart and the highest response will increase, but the position of the peak will remain the same in relation to the rest of the curve.

Both Stevens and Tsien found that the peak shifted to the right without any change in the maximum response—indicating, they argue, that the changes in LTP occur presynaptically. Stevens and postdoc John Bekkers first worked with cultured hippocampal neurons, firing individual neurons and recording hundreds of postsynaptic responses. Because there is always the possibility that cells in culture don't behave as they do in the brain, Stevens and Bekkers repeated their work in slices of whole hippocampus and found the same result—a shift in the reponse curve that could only be explained by an increase in transmitter release.

Tsien, along with Roberto Malinow, who is now at the University of Iowa, did all their experiments in brain slices and found a shift in the curve similar to that seen by Stevens. But, Tsien adds, that finding is not even their most convincing piece of evidence. The most dramatic finding is the reduction, after LTP, in the number of cases in which the presynaptic nerve cell is stimulated but no quanta are released. Such events, known as "failures," fell from an average of 63% to 17% after LTP, suggesting that there is clearly an increase in the amount of transmitter being released from the presynaptic cell. "It just pops out at you," Tsien says. "The decrease in failures is immediate, it's enormous, and it blows you away when you see it."

Even a staunch member of the opposing camp—UCSF's Roger Nicoll—concedes that "taken together, it's a very convincing story." Yet Nicoll is hardly ready to give up. The results he and Lynch obtained on receptor response during LTP still contradict Tsien's and Steven's findings, he notes, and can't be easily dismissed.

Furthermore, Nicoll and Robert Zalutsky, in this issue of *Science* (p. 1619), and Lynch and co-workers, in a recent issue of *Synapse* (vol. 5, p. 333), report some new results that would also seem to argue against the presynaptic model. Ironically, those results actually depend on the related finding that in a different type of hippocampal synapse, made by nerve fibers called mossy fibers, potentiation is clearly presynaptic. That result was obtained by showing that in the mossy fibers potentiation shares a common step with processes known to increase presynaptic transmitter release.

But in experiments that the Lynch and Nicoll teams both included as controls, they confirmed a result obtained earlier by Bruce McNaughton of the University of Colorado at Boulder: LTP at NMDA synapses shares no steps in common with known mechanisms that boost transmitter release. This finding suggests that, if LTP does depend on boosting presynaptic release, it does so through a different process than the one used under other circumstances. "It has to be by a very different mechanism," Nicoll says, "a very curious, fascinating effect."

So the debate between the presynaptic and postsynaptic camps is far from over. But at present the tide seems to have turned. Where only a year ago the evidence seemed heavily in favor of those who believed LTP could be explained by changes in the postsynaptic cell, Stevens and Tsien have tipped the balance toward the presynaptic view.

Striking as their presentation was, it is clear that the full story won't be told until considerably more supporting data are in. Among the unresolved questions if Tsien and Stevens are right is the nature of the messenger that would have to carry the LTP signal back to the presynaptic cell. Bliss has evidence for one candidate, and other labs are entering the search as well.

Having been practically alone in the presynaptic camp for years, Bliss seemed elated over Stevens' and Tsien's presentations. "It's remarkable how the seesaw has swung," he enthused. But most experienced observers of neuroscience say the seesaw ride isn't over yet. ■ MARCIA BARINAGA

Swimming for Dear Memory

Long-term potentiation (LTP), the activity-dependent strengthening of connections between nerve cells in the brain, may provide the basis for nice models of memory, but does it have anything to do with how memory actually works in an intact organism? Some evidence that it does was presented by Richard Morris of the University of Edinburgh Medical School at a recent Cold Spring Harbor symposium on the brain.

In the early 1980s, Morris devised a memory test for rats that requires them to rely on spatial memory, which is known to depend on the brain structure known as the hippocampus. The rats are placed in a pool of water clouded with milk so that they can't see below the surface. Their mission is to find a submerged platform. The reward—since rats don't like cold water—is that by climbing onto the platform they can get out of the water. Morris hypothesizes that because the platform is hidden from sight and smell, the rats must rely on spatial cues from the surroundings to remember the location of the platform.

During training sessions, normal rats learn to find the platform quickly and remember its location from one session to the next. Rats with hippocampal lesions, however, can't learn the task—they swim aimlessly and find the platform only by chance. To test the notion that this hippocampal-based learning is related to LTP, Morris made use of a pump to deliver a drug called APV into the hippocampus. APV blocks the activity of the subclass of glutamate receptors called NMDA receptors, which are required in LTP (see main story). He found that APV-treated rats were hindered in learning, but in a very specific way: although their performance improved during a training session, they couldn't remember the location of the platform from session to session as the control rats could. "What the APV is [apparently] doing is blocking the ability of the information to get into long-term memory," Morris says.

At the Cold Spring Harbor symposium, Morris presented what he believes is the strongest evidence yet that the learning defect is caused by blocking LTP. He performed dose-response studies showing that the concentration of APV in the rat's brain sufficient to block learning is "bang on top" of the concentration required to block LTP. "I think that's quite an exciting result," he says. "Going from ligand binding through to behavior with a single receptor and getting all the graphs on top of each other isn't something that happens every day."

These experiments and others by Bruce MacNaughton of the University of Colorado at Boulder are an important part of the LTP story, says Roger Nicoll, who studies LTP at the University of California, San Francisco. "They are basically the only thing we have to say that this cute cellular phenomenon has anything whatsoever to do with what it's being touted for." ■ M.B.