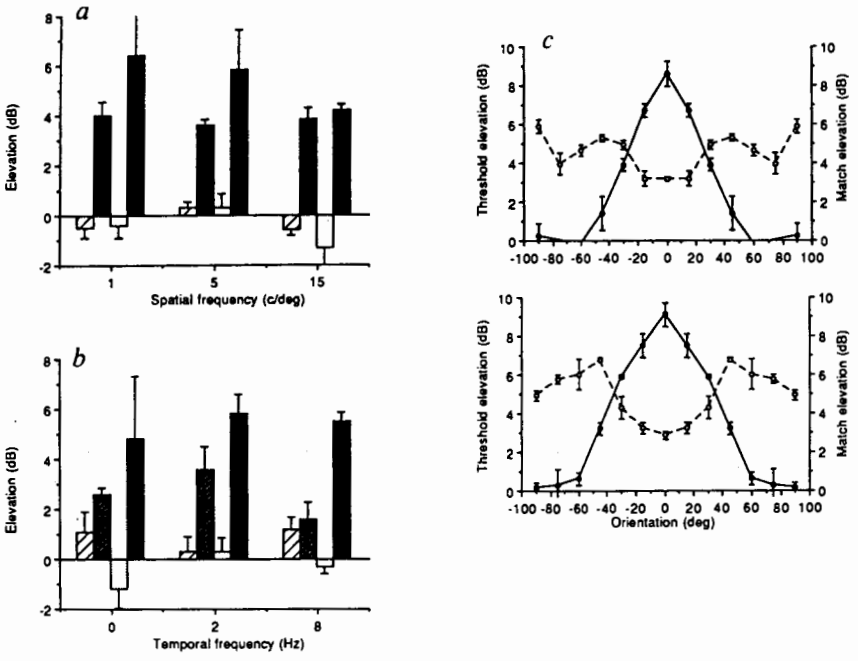

FIG. 3 *a*, The amount of threshold elevation (light bars) and the loss of contrast for a high contrast grating (dark bars) are plotted for three spatial frequencies after adapting to a horizontal pattern and testing with a vertical pattern (all patterns had a temporal frequency of 2 Hz). All other conditions are as in Fig. 1. Striped bars, subject R.S.; speckled bars, subject J.O. Error bars 1 s.e. The results are consistent in showing little or no change in threshold but a loss of suprathreshold contrast. *b*, As in *a* but for a range of temporal frequencies (spatial frequencies = 3 cycles per degree). *c*, The effects of changing the angle between the adapting and test patterns on threshold detection (solid symbols) and on the perceived suprathreshold contrast of a grating of 32% contrast (open symbols). Results for subject J.O. are shown in the upper panel and those for R.S. in the lower panel. Error bars, ±1 s.e. All conditions are as in Fig. 1, except that the adapting pattern had a contrast of 90%. To measure thresholds the same procedures were used as in Fig. 1, except that the contrast of the pattern on the unadapted side was set to 0% and subjects gave a yes/no decision on the visibility of the test pattern. This decision was used to drive the QUEST procedure.

from the adapting spatial frequency. The tuning of the spatial frequency shift across orientation is different from the classical effect[20], but corresponds well with the spatial frequency tuning of the divisive adaptation (R.J.S., manuscript in preparation).

The actual mechanisms of adaptation (in terms of the classic similar-orientation adaptation, and the new cross-orientation adaptation) have yet to be fully resolved, though one popular notion is that it involves inhibition over a protracted time course[15]. This idea is attractive as explains the results of the experiments of Greenlee and Magnussen[21]. They show that adapting to two different orientations (either simultaneously or successively) can cause less adaptation than to either one alone. This can be explained if each of the adapting patterns is inhibiting the other. Our results suggest that this will be a divisive process. The current ideas receive support from electrophysiological evidence. For example, Morrone and Burr[32] measured visually evoked potentials as a function of contrast in the presence and absence of parallel or orthogonal masking patterns. They show that parallel masks produce a shift in threshold and no change in slope, whereas orthogonal masks tend to change the slope without affecting threshold, thus complementing our results.

Cells of the striate cortex with different orientation preferences may exert inhibitory influences upon one another[9-13]. Are these inhibitory influences responsible for the present cross-orientation results? At first this might seem unlikely. If these inhbitory influences are 'adapted-out' we might expect the test pattern to appear to have a greater contrast. But if, as suggested above, adaptation reflects prolonged inhibition then the results are understandable. Recent reports of similar divisive inhibition between direction-tuned cells of area MT of the monkey[23] (see also ref. 24 for results in area 17 of the cat) might suggest that this divisive cross-dimensional inhibition is a pervasive influence throughout sensory systems. □

Received 12 August; accepted 25 October 1991.

1. Mollon, J. D. *New Scient.* **61**, 479-482 (1974).
2. Gilinsky, A. S. *J. opt. Soc. Am.* **58**, 13-17 (1968).
3. Blakemore, C. & Nachmias, J. *J. Physiol., Lond.* **213**, 157-174 (1971).
4. Blakemore, C., Muncey, J. P. J. & Ridley, R. M. *Vision Res.* **13**, 1915-1931 (1973).
5. Maffei, L., Fiorentini, A. & Bisti, S. *Science* **182**, 1036-1038 (1973).
6. Movshon, J. A. & Lennie, P. *Nature* **278**, 850-852 (1979).
7. Ohzawa, I., Sclar, G. & Freeman, R. D. *Nature* **298**, 266-268 (1982).
8. Hubel, D. & Weisel, T. *J. Physiol. Lond.* **195**, 215-243 (1968).
9. Beneveto, L. A., Creutzfelt, O. D. & Kuhnt, U. *Nature* **238**, 124-126 (1972).
10. Sillito, A. M. *Trends Neurosci.* **2**, 196-198 (1979).
11. Morrone, M. C., Burr, D. C. & Maffei, L. *Proc. R. Soc.* B**216**, 335-354 (1982).
12. Ramoa, A. S., Shadlen, M., Skottun, B. C. & Freeman, R. D. *Nature* **321**, 237-239 (1986).
13. Bonds, A. B. *Vis. Neurosci.* **2**, 41-55 (1989).
14. Georgeson, M. A. *Spat. Vision* **1**, 103-112 (1985).
15. Dealy, R. S. & Tolhurst, D. J. *J. Physiol., Lond.* **241**, 261-270 (1974).
16. Albrecht, D. G. & Hamilton, D. B. *J. Neurophys.* **48**, 217-237 (1982).
17. Tolhurst, D. J., Movshon, J. A. & Thompson, I. D. *Expl Brain Res.* **41**, 414-419 (1981).
18. Gibson, J. J. & Radner, M. *J. exp. Psychol.* **20**, 453-467 (1937).
19. Heeley, D. W. *Vision Res.* **29**, 1229-1236 (1979).
20. Blakemore, C. B. & Sutton, P. *Science* **166**, 245-247 (1971).
21. Greenlee, M. W. & Magnussen, S. *Vision Res.* **28**, 1303-1310 (1988).
22. Morrone, M. C. & Burr, D. C. *Nature* **321**, 235-237 (1986).
23. Snowden, R. J., Treue, S., Erickson, R. E. & Andersen, R. A. *J. Neurosci.* **11**, 2768-2785 (1991).
24. Dean, A. F., Hess, R. F. & Tolhurst, D. J. *J. Physiol., Lond.* **308**, 84 (1980).

ACKNOWLEDGEMENTS. We thank J. Orchard for technical assistance.

Tit for tat in heterogeneous populations

Martin A. Nowak* & Karl Sigmund†

* Department of Zoology, University of Oxford, South Parks Road, Oxford OX1 3PS, UK
† Institut für Mathematik, Universität Wien, Strudlhofg. 4, A-1090 Wien, Austria

THE 'iterated prisoner's dilemma' is now the orthodox paradigm for the evolution of cooperation among selfish individuals. This viewpoint is strongly supported by Axelrod's computer tournaments, where 'tit for tat' (TFT) finished first[1]. This has stimulated interest in the role of reciprocity in biological societies[1-8]. Most theoretical investigations, however, assumed homogeneous populations (the setting for evolutionarily stable strategies[9,10]) and programs immune to errors. Here we try to come closer to the biological situation by following a program[6] that takes stochasticities into account and investigates representative samples. We find that a small fraction of TFT players is essential for the emergence of reciprocation in a heterogeneous population, but only paves the way for a more generous strategy. TFT is the pivot, rather than the aim, of an evolution towards cooperation.

The simple 'prisoner's dilemma' is a game with two players, each having two options, C (to cooperate) and D (to defect). We use Axelrod's payoff values: if both players cooperate, both obtain $R = 3$ points; if both defect, each receives $P = 1$ point; if one player defects and the other cooperates, the defector gets $T = 5$ points and the cooperator $S = 0$. Clearly strategy D is

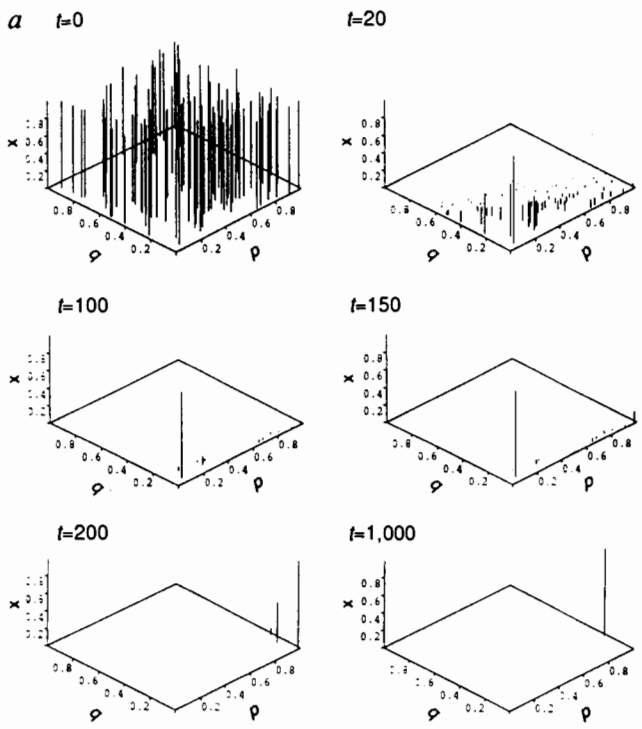

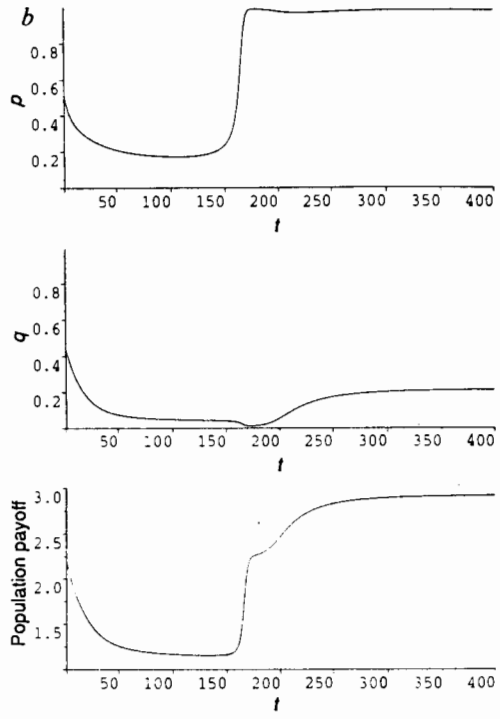

FIG. 1 The role of TFT for the emergence of cooperation in heterogeneous ensembles of strategies is illustrated in this simulation. The population consists of a randomly chosen sample of 99 reactive strategies $E_i=(p_i, q_i)$ uniformly distributed on the unit square. The stochastic TFT-like strategy $E=(0.99, 0.01)$ is then added to the population. The payoff for a strategy E_i against E_j in the infinitely IPD is obtained as $A(E_i, E_j)=1+4c'-c-cc'$, with $c=[q_i+(p_i-q_i)q_j]/[1-(p_i-q_i)(p_j-q_j)]$ and c' the analogous expression with i and j interchanged. Initially all strategies are present in the same frequency, 1%. If x_i denotes the frequency of E_i in one generation, then its frequency in the next generation will be given by $x_i'=x_i f_i(x)/\bar{f}$ for $i=1,\ldots,100$. Here $f_i(x)$ denotes the average payoff for E_i in a population given by $x=(x_1 \cdots x_{100})$ and $\bar{f}=\sum x_i f_i(x)$ is the average payoff in the population. During the first 100 generations the strategies close to AllD dominate the population. After about 150 generations stochastic TFT reappears and takes over. Reciprocity is established. This enables more forgiving strategies to grow. The final winner is GTFT ($p \simeq 0.99$, $q \simeq 0.33$). *a*, Relative frequencies (normalized such that the relative frequency of the most abundant strategy equals one) for all strategies after 0, 20, 100, 150, 200 and 1,000 generations. In our plot, strategies with very low frequencies are no longer visible, although they are still present in the numerical computations. *b*, Population averages for *p*, *q* and payoff. The apparition of stochastic TFT increases average payoff from about 1.2 to 2.25 (the payoff for the completely random strategy). To catalyse the turn of the tide, a strategy has to be very close to TFT. A strategy like (0.9, 0.1) will not, in general, be close enough to serve this purpose. Fully cooperative play (payoff 3) is only established after GTFT has taken over. Generosity pays under conditions of uncertainty.

best, whatever the other player does. Hence both players will use D and end up with one point only, instead of three points for mutual cooperation. This has nothing to do with rationality or foresight. We have in mind biological applications, where the payoff is number of offspring (more successful strategies have higher fitness) and strategies are inherited. The defectors, then, will necessarily out-compete the cooperators.

For the 'iterated prisoner's dilemma' (IPD), one assumes a constant probability w for another round; w can also be viewed as discount factor for future payoff. A strategy is a stochastic or deterministic rule specifying the choice of C or D in every round as a function of the history of the interaction so far.

In Axelrod's round-robin tournaments, TFT did best. It consists in playing C in the first round and from then on doing whatever the other player did in the previous round. Axelrod stressed, however, that TFT is not the best strategy (for $w > (T-R)/(T-P)=2/4$, there is no best strategy in the IPD), and that in a pairwise interaction with any given partner, TFT never does better than this partner. But TFT can elicit cooperation from players who might defect among themselves.

There are essentially two ways of studying TFT's evolutionary role. One can check whether a homogeneous TFT population will resist invasion by mutant strategies. No single mutant can do better than TFT[1] if $w > \max((T-R)/(T-P), (T-R)/(R-S)=2/3$. But AllC (which always cooperates) and many other strategies do as well as TFT.

The heteromorphic approach starts with a widely scattered distribution and subjects it to selection. This has been done in Axelrod's 'ecological' kind of tournament[1], when 63 entries were submitted, generation after generation, to a round-robin contest, their frequencies proportional to their payoff in the previous generation. The main problem here is to find plausible values from which to start. The participants in Axelrod's tournament included a good sample of game-theorists, but are unlikely to be representative of simpler biological communities.

The array of strategies for the IPD is so huge that it cannot be sampled by statistical means and has to be drastically reduced. Axelrod[11], for instance, considered deterministic strategies where the decision to cooperate or defect in each round depended on the outcome of the three previous rounds. Here, we consider a situation both simpler and more complex: the decision depends only on the previous round, but is stochastic and not deterministic. It reflects the tendency to cooperate in answer to the other's move. This is meant to apply to biological interactions where memories are short and decisions uncertain. Players may misinterpret the other's move (or identity), or misimplement their own intention. This relates not only to the 'trembling hand' behind the concept of perfect equilibrium[12,13], but also to the 'blurred minds'.

Errors are costly to a TFT population[1]. Mistakes gives rise to sequences of alternating defections. This weakness of TFT is most pronounced in interactions against itself. When Axelrod repeated his tournament with an error rate of 1%, TFT still finished first, because the field of contestants was so diversified, but in a monomorphic TFT-population, the smallest error rate reduces fitness by 25%. It may be better to forgive with

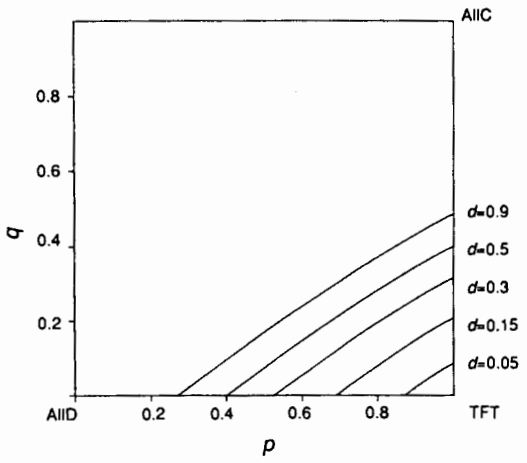

FIG. 2 Strategies able to invade AllD in a cluster. For given cluster size d, when can an AllD population be invaded by a cluster of (p, q) strategists? If d is the frequency of (p, q) and $1-d$ that of AllD, the frequency of (p, q) grows if and only if $d > (1-p+q)^2/[3p-4q-3(p-q)^2]$. For a given d, this defines a neighbourhood of TFT in the unit square of strategies, drawn here for $d = 0.9$, 0.5, 0.3, 0.15 and 0.05. If d decreases, the neighbourhood shrinks to the corner point (1, 0). Thus TFT is the strategy that can invade defectors in a minimal cluster.

a certain probability.

A reactive strategy is given by a triple (y, p, q), where y is the probability to cooperate in the first round and p and q are the conditional probabilities to cooperate after a C (respectively D) of the other player[14-16]. TFT corresponds to (1, 1, 0), AllD to (0, 0, 0), AllC to (1, 1, 1), 'suspicious tit-for-tat' (STFT) to (0, 1, 0), and so on. We shall mostly be interested in properly stochastic strategies, with $0 < y, p, q < 1$.

For simplicity, we consider only the infinitely IPD (the limit case, $w = 1$). The initial move has no role, then, as its effect is 'forgotten' in the long run. A strategy now will simply be a point (p, q) in the unit square. The assumption $w = 1$ is only to keep things simple. The results still hold if we allow for some discount of the future ($w < 1$).

A simple characterization[16] of when a strategy can be invaded by another strategy shows that 'generous tit-for-tat' (GTFT), the strategy $(1, q)$ with $q = \min\{1-(T-R)/(R-S), (R-P)/(T-P)\} = \frac{1}{3}$, is optimal in the sense that among all reactive strategies immune to invasion by less cooperative strategies (lower p or q values), it affords the highest payoff for a population adopting it[15,17].

If we choose a representative sample of reactive strategies E_1 to E_n, initially all equally frequent, we can watch their frequencies evolve under the action of selection. With $n = 100$ different reactive strategies uniformly distributed on the unit square, evolution proceeds in most cases towards AllD: those (p, q)-strategies from the sample which are closest to (0, 0) increase in frequency, while all others vanish. This follows because a large percentage of the random sample has high q values and does not retaliate against exploiters. With a rich diet of 'suckers', it pays to defect.

The outcome alters dramatically if one of the initial strategies (added by hand or by chance), is TFT, or very close to it (Fig. 1). The first phase is practically indistinguishable from the previous run. The strategies near AllD grow rapidly. TFT and all other reciprocating strategies (near (1, 0)) seem to have disappeared. But an embattled minority remains and fights back. The tide turns when 'suckers' are so decimated that exploiters can no longer feed on them. Slowly at first, but gathering momentum, the reciprocators come back, and the exploiters' now wane. But the TFT-like strategy that caused this reversal of fortune is not going to profit from it: having eliminated the exploiters, it is robbed of its mission and superseded by the strategy closest to GTFT. Evolution then stops. Even if we introduce occasionally 1% of another strategy, it will vanish.

Are reactive strategies a representative ensemble? They cover a broad range of cooperative, defective, fully random or reciprocal behaviour. They include strategies like AllC, AllD, TFT and GTFT. Their stochasticity reflects life's fuzziness. Their simplicity seems to match the metaphorical nature of the 'prisoner's dilemma'. That simple strategies are important for the IPD is shown by TFT's triumph in Axelrod's tournaments.

Will these findings hold for larger sets of strategies? Our simulations can be extended to more complex strategies and $w < 1$. Usually, however, the strategy space becomes so large that representative samples are too unwieldy. We therefore sketch only one extension, which covers stochastic strategies depending on the other player's last move, but also one's own. They are given by the four conditional probabilities p_1, p_2, p_3 and p_4 to cooperate after (C, C), (C, D), (D, C) [respectively (D, D)] in the previous round. In these simulations behaviour is similar: strategies near TFT ($p_1 = p_3 = 1$; $p_2 = p_4 = 0$) favour the emergence of strategies near GTFT ($p_1 = p_3 = 1$; $p_2 = p_4 = 1/3$).

Axelrod[11] simulated the evolution of deterministic strategies depending on the previous three moves by both players (a promising variant including noise is described in ref. 18). Together with the specifications for the first rounds, such strategies are represented by strings of 70 Cs and Ds. Axelrod used a genetic algorithm, starting with a random sample of strategies, applying game dynamics and occasionally modifying strategies by mutating letters or recombining strings. Evolution always started off towards defection, but then veered towards cooperation. One can ask whether such a change is always due to the occurrence of a TFT-like strategy. Stochastic strategies like GTFT cannot occur here: but it would be interesting to know whether, here again, a generous strategy (for example, tit for two tats) would be the ultimate beneficiary of the action of stern retaliators. Or does forgiveness pay only under uncertainty? Axelrod and Dion[4] suggest that for larger noise "generosity invites exploitation". In fact, the optimal forgiveness in reactive strategies decreases if noise grows[15].

In our simulations, TFT is almost specified by its police role: strategies that are not very close to TFT do not have its effect. But an evolution twisted away from defection (and hence due to TFT) leads not to the prevalence of TFT, but towards more generosity. TFT's strictness is salutary for the community, but harms its own.

TFT acts as a catalyser. It is essential for starting the reaction towards cooperation. It needs to be present, initially, only in a tiny amount; in the intermediate phase, its concentration is high; but in the end, only a trace remains. □

Received 19 August; accepted 5 November 1991.

1. Axelrod, R. *The Evolution of Cooperation* (Basic Books, New York, 1984).
2. Axelrod, R. & Hamilton, W. D. *Science* **211**, 1390-1396 (1981).
3. Trivers, R. *Social Evolution* (Cummings, Menlo Park, 1985).
4. Axelrod, R. & Dion, D. *Science* **242**, 1385-1390 (1988).
5. Boyd, R. & Lorberbaum, J. P. *Nature* **327**, 58-59 (1987).
6. May, R. M. *Nature* **327**, 15-17 (1987).
7. Milinski, M. *Nature* **325**, 434-435 (1987).
8. Wilkinson, G. S. *Nature* **308**, 181-184 (1984).
9. Maynard Smith, J. *Evolution and the Theory of Games* (Cambridge University Press, UK, 1982).
10. Hofbauer, J. & Sigmund, K. *The Theory of Evolution and Dynamical Systems* (Cambridge University Press, UK, 1988).
11. Axelrod, R. in *Genetic Algorithms and Simulated Annealing* (ed. Davis, D.) (Pitman, London, 1987).
12. Selten, R. *Internat. J. Game Th.* **4**, 25-55 (1975).

13. Boyd, R. *J. theor. Biol.* **136**, 47-56 (1989).
14. Nowak, M. & Sigmund, K. *J. theor. Biol.* **137**, 21-26 (1989).
15. Nowak, M. & Sigmund, K. *Acta appl. Math.* **20**, 247-265 (1990).
16. Nowak, M. *Theor. Pop. Biol.* **38**, 93-112 (1990).
17. Molander, P. *J. Conflict Resolut.* **29**, 611-618 (1985).
18. Lindgren, K. in *Artificial Life II* (eds Farmer, D. *et al.*) (Proc. Santa Fe Inst. Stud., Addison-Welsey, 1991).

ACKNOWLEDGEMENTS. Support from the Austrian Forschungsförderungsfond is gratefully acknowledged.

Regulation of the G1–S transition in postembryonic neuronal precursors by axon ingrowth

Scott B. Selleck, Cayetano Gonzalez*, David M. Glover* & Kalpana White

Department of Biology, Brandeis University, Waltham, Massachusetts 02254-9110, USA
* Cancer Research Campaign Laboratories, Cell Cycle Genetics Group, Department of Biochemistry, The University, Dundee DD1 4HN, UK

IN the newly cellularized *Drosophila* embryo, progress through the cell cycle is regulated at the G2–M transition[1,2]. We have examined cell-cycle regulation later in *Drosophila* development, in a group of postembryonic neuronal precursors. The S-phase precursor cells, which generate photoreceptor target neurons (lamina neurons) in the central nervous system, are not present in the absence of photoreceptor innervation[3]. Here we report that axons selectively approach G1-phase precursors. Without axon ingrowth, lamina precursors do not enter their final S phase and by several criteria, arrest in the preceding G1 phase. These findings provide evidence that at this stage in development the control of cell division can occur at the G1–S transition.

Neurons of the adult optic lobe arise from cell divisions during postembryonic life[4]. Lamina neurons, which receive direct synaptic input from photoreceptors, are derived from a discrete group of S-phase lamina precursor cells (LPCs)[3,5]. S-phase LPCs are part of a proliferative epithelium on the surface of the brain, the outer proliferative centre (OPC). The OPC epithelium invaginates anterior to the developing lamina to form a 'V'-shaped furrow (Fig. 1*b*, *d*). Pulse-labelling with bromodeoxyuridine (BUdR) shows three OPC S-phase domains. They are (from anterior to posterior): (1) the anterior OPC (aOPC), a broad belt of cells; (2) scattered BUdR-incorporating cells slightly more posterior; and (3) S-phase LPCs (Fig. 1*a*, *b*). To determine the origins and fate of S-phase LPCs, larvae were pulsed with BUdR and the labelled cells followed after varying lengths of time. Six hours after BUdR injection, cells labelled as S-phase LPCs (red cells in Fig. 1*d*, *f*, *h*) have moved into the lamina *en bloc* (compare Fig. 1*c*, *d* with *e*, *f*). By 12 h after BUdR delivery, labelled cells originating from the anterior region of the OPC (blue cells, Fig. 1*d*, *f*, *h*) have travelled along the furrow to reside where S-phase LPCs were seen with a pulse label (compare Fig. 1*c*, *d* with *g*, *h*). (aOPC progeny also become

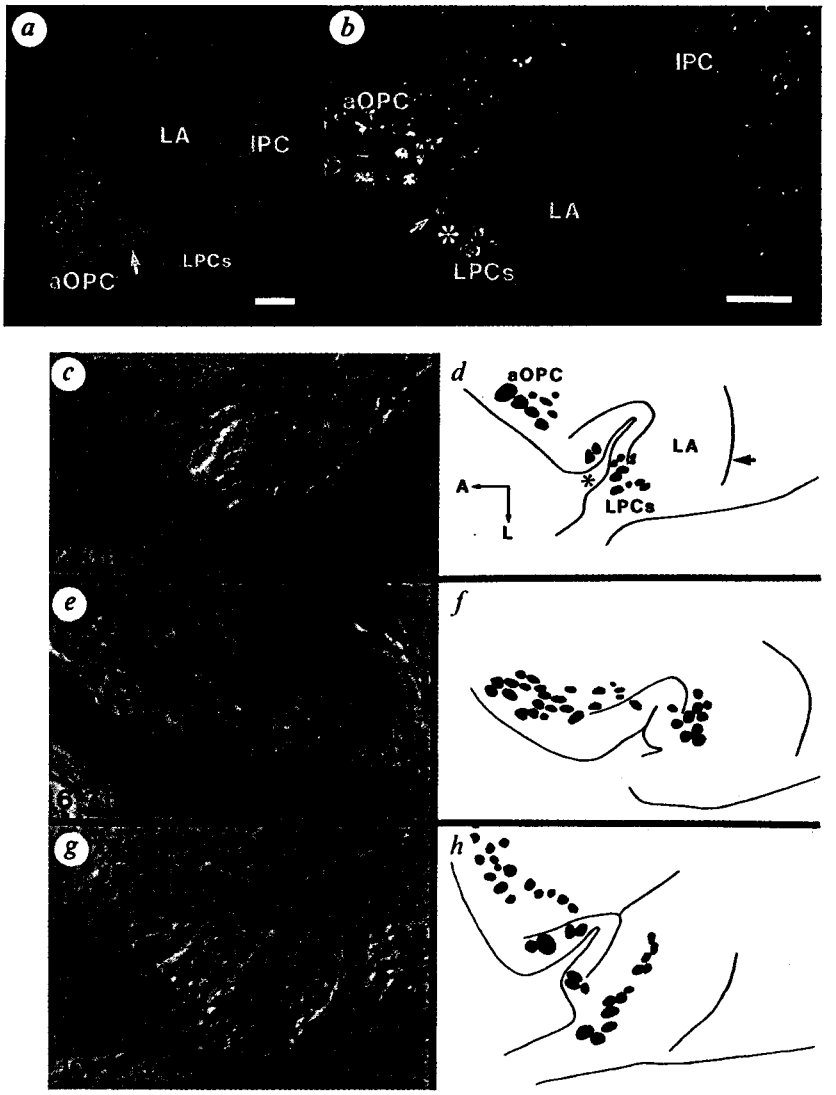

FIG. 1 Anatomy of the outer proliferative centre, and origins of S-phase LPCs. *a*, *b*, S-phase cells in the third instar larval CNS visualized by Bromouridine deoxyribose (BUdR) incorporation (shown in green). Propidium iodide-stained nuclei shown in grey tones. Confocal images showing lateral (*a*, anterior left, ventral down) and horizontal (*b*, anterior left, lateral down) views. The OPC S-phase domains are: (1) the anterior OPC (aOPC); (2) scattered cells along the anterior segment of the furrow (white arrows *a*, *b*); and (3) S-phase LPCs. Asterisk indicates furrow; black arrows, posterior limit of the lamina (LA); IPC, inner proliferative centre. Scale bars, *a*: 25 μm, *b*: 10 μm. *c*–*h*, BUdR pulse-chase analysis of LPCs. Cells pulse-labelled with BUdR (*c*, *d*, 2–3 h post-BUdR injection) were followed at 6–7 h (*e*, *f*) and 10.5–11.5 h (*g*, *h*) post-labelling. *c*, *e*, *g*, Paraffin sections cut in horizontal plane (anterior left, lateral down). BUdR-labelled cells show brown staining. *d*, *f*, *h*, Representations of sections shown in *c*, *e*, *g*, respectively. Only cells that incorporate BUdR while part of the OPC, and OPC progeny are indicated. Cells labelled while in the S-phase domains anterior to S-phase LPCs are represented in blue, cells labelled as S-phase LPCs in red. BUdR incorporating cells of the IPC, the occasional cell in the lamina and perineural sheath are not indicated in the diagrams. Abbreviations and markings as for *a*, *b*. Scale bar, 10 μm.

METHODS. Wild-type (Canton-S) CNS preparations shown in *a*, *b* were labelled for 0.5 h in Schneider's medium containing BUdR (ref. 14), and viewed with a Biorad MRC 600 confocal microscope. For the pulse-chase analysis, wild-type climbing third instar larvae were injected with BUdR (refs 3, 14) and incubated at 25 °C before dissection. BUdR-incorporating cells were detected using anti-BUdR antibody (Becton-Dickinson)[3,14] and 6 μm paraffin sections made[15].