THE ROYAL SOCIETY

Analysis of cytokine dynamics in corneal allograft rejection

Cliburn C. W. Chan[1*†], **Jaroslav Stark**[1†] **and Andrew J. T. George**[2†]

[1] *Centre for Nonlinear Dynamics and its Applications, University College London, Gower Street, London WC1E 6BT, UK*
[2] *Department of Immunology, Division of Medicine, Imperial College School of Medicine, Hammersmith Hospital, Du Cane Road, London W12 0NN, UK*

Motivated by the discovery of oscillations in tumour necrosis factor-α (TNF-α) concentration in the aqueous humour of rabbits undergoing corneal allograft rejection, a simple mathematical model is developed for the regulation of TNF-α, which incorporates both negative feedback and amplification pathways. Mathematical analysis reveals a surprisingly rich behavioural repertoire for this simple cytokine pathway, including excitability, threshold behaviour, hysteresis, oscillations and bistability. This suggests new possibilities for experimental demonstration, and reveals the potential contributions of nonlinear dynamics to understanding cytokine regulation.

Keywords: tumour necrosis factor-α; corneal transplantation; oscillations; nonlinear dynamics; mathematical model; cytokine network

1. INTRODUCTION

The cornea is a classical site of immune privilege, with limited major histocompatibility complex (MHC) expression, constitutive expression of Fas ligand, and local release of immunosuppressive cytokines and neuropeptides in addition to the sequestration of antigen by the blood–ocular barrier (Niederkorn 1990; Streilein *et al*. 1992; Ferguson *et al*. 1995; Griffith *et al*. 1995; Williams & Coster 1997). However, corneal grafts show a significant failure rate (about 25% at four to five years) (Williams *et al*. 1992), the major cause of which is immunological rejection, emphazising that immune privilege is a relative concept.

The immune events culminating in graft rejection are orchestrated by a complex network of cytokines, the regulation of which is still poorly understood. One of the major cytokines involved is tumour necrosis factor-α (TNF-α), which is critical in the successful initiation, maintenance and resolution of inflammation (Strieter *et al*. 1993).

The regulation of TNF-α is complex. TNF-α has the potential for amplifying its own effects, both directly by autocrine and paracrine pathways (Philip & Epstein 1986; Amiot *et al*. 1997), and indirectly by enhancing chemotaxis, adhesion and transmigration of mononuclear cells (Tracey & Cerami 1994). In addition, several cytokines including interferon-α(IFN-α), interferon-γ (IFN-γ), interleukin-1 (IL-1), IL-4, IL-10, IL-13, and transforming growth factor-β (TGF-β), many of which are induced by TNF-α, can up- and downregulate the expression of TNF-α. Finally, the shedding of soluble TNF receptors (sTNFR) stimulated by TNF-α itself can also contribute to the damping of TNF-α effects (Heaney & Golde 1996).

We have recently discovered that TNF-α concentrations in the aqueous humour from rabbits with corneal allografts undergo striking oscillations, which are not observed in control rabbits with corneal autografts. These sustained oscillations have large amplitudes, although we are unable to determine if they are regular. These data are described in detail elsewhere (Rayner *et al*. 1999). A representative example of the data from one animal is shown in figure 1.

We postulated that these oscillations arose from the regulatory interactions between TNF-α and its inhibitors, the most probable candidates being IL-10, TGF-β and sTNFR.

Oscillations often reveal interesting underlying dynamics, which are best approached via a formal mathematical model. We therefore attempted to understand the mechanism behind the TNF-α oscillations observed in rabbit corneal allograft by constructing a simple model of TNF-α regulation in the eye.

The model predicts that such a cytokine network would show a rich set of behaviours under certain quantifiable conditions, including excitability, oscillations, existence of a threshold and bistability.

2. PRODUCTION OF A MATHEMATICAL MODEL FOR THE REGULATION OF TNF-α IN CORNEAL ALLOGRAFT REJECTION

The essential features of the model are illustrated in figure 2. The various inhibitors of TNF-α have been consolidated into a single generic inhibitor, and the known positive and negative feedback pathways shown.

This informal pictorial model is converted into a system of ordinary differential equations by simply

*Author for correspondence.
† Secondary affiliation of all authors: CoMPLEX, Centre for Mathematics and Physical Sciences in the Life Sciences and EXperimental Biology, University College London, Gower Street, London WC1E 6BT, UK.

Received 24 May 1999 *Accepted* 2 August 1999

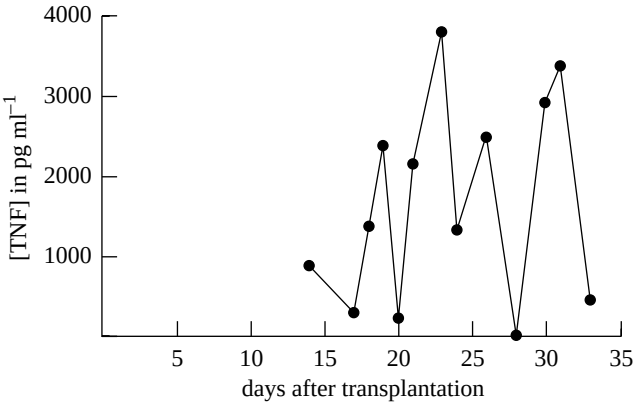

Figure 1. Representative example of experimentally observed TNF-α fluctuations in the aqueous humour of a rabbit undergoing corneal graft rejection. Data from one animal out of ten. The levels of TNF-α were measured using a bioassay in which the death of actinomycin-sensitized L929 cells in response to TNF-α was measured. Experiments were carried out in triplicate.

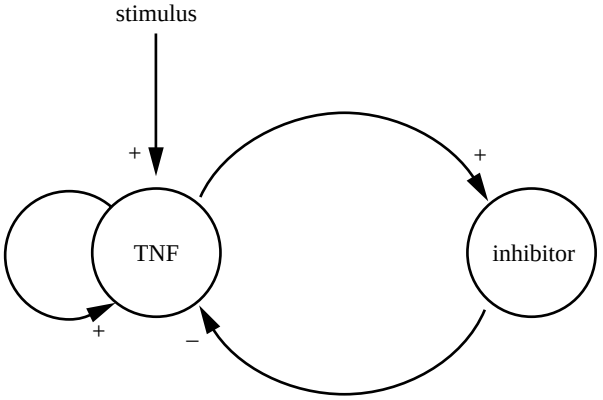

Figure 2. Model for the cross-regulation of TNF-α and inhibitor.

assuming that the rate of change of TNF-α and inhibitor concentration is the difference between its production and clearance.

3. DETAILS OF THE MATHEMATICAL MODEL

A first-order differential equation is written for each cytokine involved as a simple balance between production and clearance, i.e.

rate of change of cytokine concentration
= rate of cytokine production − rate of cytokine clearance

This leads to the following coupled ordinary differential equations for the model:

$$\frac{\mathrm{d}x}{\mathrm{d}t} = v_1 \frac{x^n + \epsilon_1^n}{x^n + \alpha^n} \frac{\beta}{y + \beta} - d_1 x,$$
$$\frac{\mathrm{d}y}{\mathrm{d}t} = k_2 + v_2 \frac{x + \epsilon_2}{x + \gamma} - d_2 y, \qquad (1)$$

where the biological interpretation of the various variables are as follows: x, concentration of TNF-α; y, concentration of inhibitor; v_1, maximal velocity of TNF-α production (corresponds to stimulation, e.g. by antigen); v_2, maximal velocity of inhibitor production; k_2, constant rate of inhibitor production which is not feedback regulated; d_1, clearance rate for TNF-α; d_2, clearance rate for inhibitor; n, Hill coefficient which determines the steepness of the dose–response curve; α, threshold for TNF-α positive feedback on rate of TNF-α production; β, threshold for inhibitor negative feedback on rate of TNF-α production; γ, threshold for TNF-α positive feedback on rate of inhibitor production; ϵ_1, baseline for the function representing TNF-α positive feedback on rate of TNF-α production, $\epsilon_1 \ll \alpha$; ϵ_2, baseline for the function representing TNF-α positive feedback on rate of inhibitor production, $\epsilon_2 \ll \gamma$.

The maximal rate of TNF-α synthesis is assumed to be a function of the stimulation to the system.

This might be due to mitogens or to the level of antigen in the system (acting for example via cytokine release by T cells or engagement of antibody coated material). Thus,

$$v_1 = f(s), \qquad (2)$$

where s is the degree of stimulation and $f(s)$ is an increasing saturable function.

The feedback loops are modelled using Hill functions representing dose–response curves. Hill functions can be derived from simple mass action kinetics (Segel 1984), and have the general form $h(x) = x^n/(s^n + x^n)$ where x represents cytokine concentration, s is the concentration where the reaction velocity is 50% of its maximum possible, and n is a parameter (known as the Hill coefficient) determining the slope of the function. In particular, the positive feedback function of TNF-α on its own production is assumed to be sigmoidal, because recruitment of mononuclear cells is likely to show a nonlinear dependence on the concentration of TNF-α. It is not critical whether the other feedback functions have sigmoidal or hyperbolic dose–response curves because qualitatively similar behaviour can be observed in numerical simulations. The clearance of cytokine and inhibitor is assumed to be linear.

By a suitable choice of change of variables, we can arrive at a system of equations in which the variables are dimensionless. Using this technique of non-dimensionalization, an equivalent set of equations with fewer parameters can be derived by substituting $u = x/\alpha$, $v = y/\beta$, and $\tau = d_1 t$,

$$f(u,v) = \frac{\mathrm{d}u}{\mathrm{d}\tau} = A \frac{u^n + E_1^n}{u^n + 1} \frac{1}{v + 1} - u,$$
$$g(u,v) = \frac{\mathrm{d}v}{\mathrm{d}\tau} = B + C \frac{u + E_2}{u + F} - Dv, \qquad (3)$$

with the parameters $A = v_1/d_1\alpha$, $B = k_2/d_1\beta$, $C = v_2/d_1\beta$, $D = d_2/d_1$, $E_1 = \epsilon_1/\alpha$, $E_2 = \epsilon_2/\alpha$, and $F = \gamma/\alpha$.

The above model as outlined in figure 2 will not be new to immunologists because similar informal representations of cytokine networks are commonly used. In formalizing such network models in mathematical terms, we are simply stating our assumptions and biases explicitly in the form of ordinary differential equations, in order that the behaviour of the system can be studied using the tools of nonlinear dynamics.

Our first test of the model is to confirm if it can replicate previously confirmed experimental results, and provide a possible explanation for the observed oscillatory behaviour. We then use bifurcation theory to characterize the qualitatively different solutions of the model, in order to understand the range of physiological behaviours potentially available to the network. Interestingly, we find solutions that predict unusual responses not previously documented in the literature, which offer new experimental tests of cytokine network regulatory mechanisms.

Clinicians too may find the counter-intuitive behaviour of aspects of the model a useful caution, as they start using cytokine agonists and antagonists therapeutically to perturb the cytokine network. A better formal understanding of the regulation of cytokine networks may help elucidate the paradoxes of success and failure observed in clinical trials of cytokine-related therapy.

4. RESULTS

Results from the model come from a combination of mathematical analysis and numerical computation. Only the main results of biological interest will be presented here.

(a) *Effect of a single antigen bolus*

We have simulated the effect of injecting a single antigen bolus by assuming that the antigen concentration decreases exponentially with time, and letting the stimulation parameter A be related to antigen concentration s by a hyperbolic function, i.e.

$$\frac{\mathrm{d}s}{\mathrm{d}t} = -d_3, \qquad A = \frac{s}{k_3 + s}, \tag{4}$$

where d_3 and k_3 are constants. There is an initial surge in the concentration of TNF-α in response to the stimulation produced by the antigen, which then decreases gradually to negligible levels as the antigen is cleared. This is consistent with experimental results, for example, the rapid induction of TNF-α after injection of tumour necrosis producing (TNP)-spleen cells into the anterior chamber (Ferguson *et al*. 1994).

This is an obvious result, and simply shows that the prediction of the mathematical model agrees with intuition.

(b) *Effect of sustained stimulation to the system*

The maximal rate of TNF-α production reflects the stimulation via antigen or mitogen, because it is a function of only the antigen concentration. We can therefore simulate different antigen loads by increasing this parameter. In the context of transplantation we assume that the stimulation rapidly rises to a peak post grafting, remaining at a plateau and changing only slowly during the active phase of graft rejection, and finally decreasing when the graft is 'burnt-out' and destroyed. The assumption of constant antigen stimulation is likely to be true to a first approximation also for autoimmune disease, chronic graft-versus-host disease and chronic infections.

We can therefore simulate the increase in stimulation by increasing the parameter corresponding to the maximal rate of TNF-α production. Throughout the simulation, the antigen load is assumed to be constant, because we do not consider the final phase of rejection.

There are five main qualitatively different phase portraits for different antigen loads, depending on the other fixed parameters.

Type I: single sink. There is a unique equilibrium concentration of TNF-α and inhibitor, and all initial conditions converge to this.
Type II: two sinks. There are two stable equilibrium concentrations of TNF-α and inhibitor. The system converges to one or the other depending on the initial conditions.
Type III: stable cycle. There is a stable periodic oscillation of TNF-α and inhibitor, and all initial conditions converge to this.
Type IV: sink and stable cycle. There is a stable equilibrium together with a periodic oscillation of TNF-α and inhibitor. The system converges to either an equilibrium state or an oscillatory state depending on the initial conditions.
Type V: excitable behaviour. There is a unique stable equilibrium concentration of TNF-α and inhibitor, but the system may converge to this via either a short route or a long excursion, depending on initial conditions.

These five phase portraits together with their 'typical' time-series are shown in figure 3.

(i) *Type I behaviour: single sink*

Here there is only a single sink, and the concentration of TNF-α will always eventually settle to this equilibrium value given sufficient time. The concentration of inhibitor will also come to a steady state. Such behaviour would correspond to the expected behaviour of the system, where the level of cytokine rises in response to the stimulus, reaches equilibrium and remains at a steady state for as long as the stimulus remains (figure 3*a*).

(ii) *Type II behaviour: two sinks*

Depending on the initial TNF-α and inhibitor concentrations, the final value of TNF-α can be either of the two equilibrium concentrations, as shown in the time-series. Similarly, there are two corresponding equilibrium values for the inhibitor concentration. The two sinks are separated by a saddle. While the levels of TNF-α at either of the sinks will be in stable equilibrium, it is possible to switch from one sink to the other by adding (removing) sufficient TNF-α or inhibitor to (from) the system. This is an example of bistability, where two stable solutions coexist at the same parameter values.

In biological terms, this would mean that under the same conditions, the system could have relatively high or low TNF-α levels which would be stable unless there was a sufficient perturbation of the system, whereupon there would be a sudden change to the other state. Which equilibrium the system settles to would be a function of the initial starting conditions.

(iii) *Type III behaviour: stable cycle*

The TNF-α level never settles to an equilibrium, but instead converges to an oscillation with constant amplitude and period. The inhibitor concentration will also

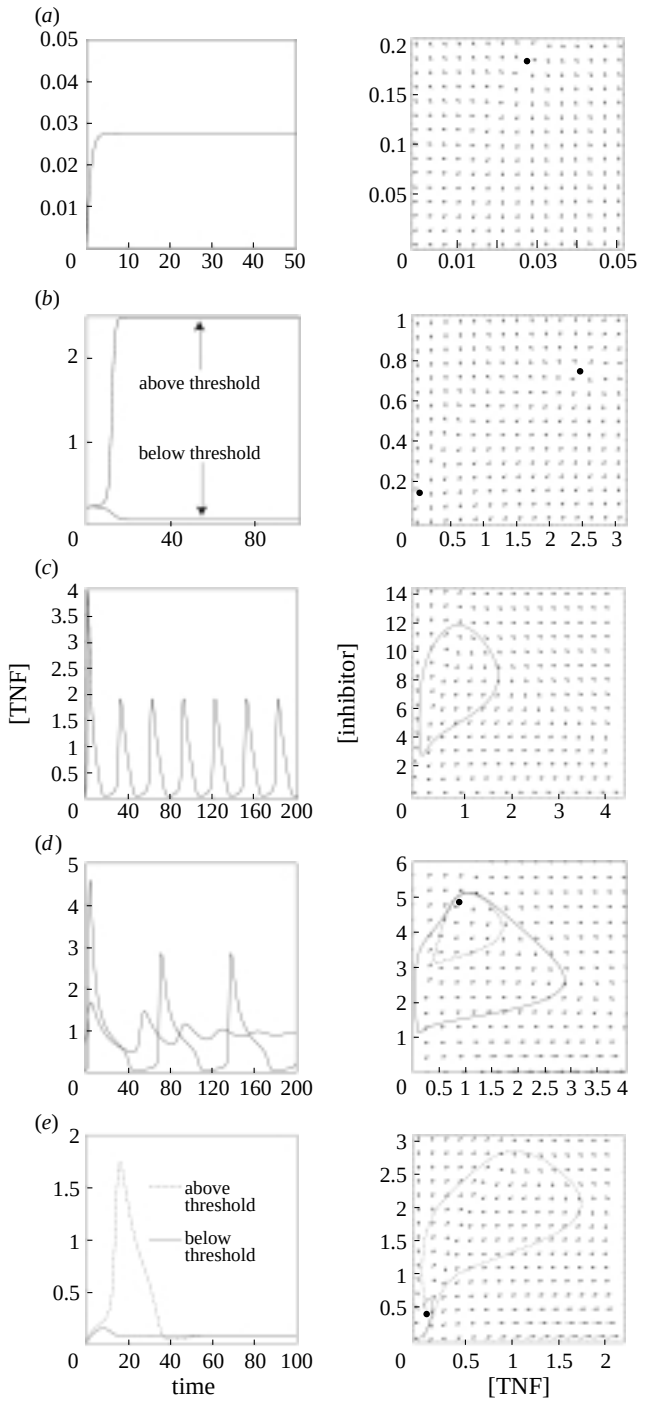

Figure 3. The figure shows type I, II, III, IV and V behaviour (*a–e*, respectively). (*a*) Single sink: stable steady behaviour where $A = 3$, $B = 0$, $C = 5$, $D = 1$, $E_1 = 0.1$, $E_2 = 0.01$, $F = 1$, $n = 2$; (*b*) two sinks: threshold behaviour where $A = 5$, $B = 0$, $C = 1$, $D = 1$, $E_1 = 0.1$, $E_2 = 0.1$, $F = 1$, $n = 2$; (*c*) stable cycle: sustained oscillation where $A = 21$, $B = 0$, $C = 5$, $D = 0.2$, $E_1 = 0.1$, $E_2 = 0.01$, $F = 1$, $n = 2$; (*d*) sink and stable cycle: coexistence of oscillations and steady state where $A = 11.6$, $B = 0$, $C = 1$, $D = 0.1$, $E_1 = 0.1$, $E_2 = 0.01$, $F = 1$, $n = 2$; (*e*) excitable behaviour where $A = 7$, $B = 0$, $C = 1$, $D = 0.18$, $E_1 = 0.1$, $E_2 = 0.01$, $F = 1$, $n = 2$. The time-series is shown on the left-hand side and the corresponding phase portrait–vector field on the right-hand side. The antigenic stimulation is assumed to be constant. See the text for a description of behaviour. Filled circles, sinks; solid closed loops, stable cycles; dashed closed loops, unstable cycles.

oscillate, lagging behind the TNF-α oscillations (not shown). This may be similar to the fluctuations in TNF-α levels shown in our transplant model.

(iv) *Type IV behaviour: sink and stable cycle*

Depending on the initial TNF-α and inhibitor concentrations, the final result may be either an equilibrium or an oscillation. This is therefore a different form of bistability. If such a system could be observed experimentally, it would be possible to stop oscillating TNF-α and inhibitor concentrations by simply adding or removing (e.g. with an antibody) TNF-α or inhibitor. Conversely, a steady state concentration of TNF-α and inhibitor can similarly be induced to oscillate by adding or removing TNF-α or inhibitor. However, it is unlikely such bistability is biologically relevant in the corneal allograft model for two reasons. (i) The basin of attraction of the sink is very small compared with that of the cycle, in all the numerical simulations we have run. This means that most combinations of TNF-α and inhibitor concentration found in the initial state will eventually end up on the cycle and thus show oscillatory behaviour, while only a small fraction of combinations of TNF-α and inhibitor initial concentrations lead to the sink. (ii) More importantly, the region of parameter space (see below) where bistability exists is very small. This means that the parameter values necessary for such behaviour are very limiting. With the inevitable noise and variation in biological systems, it is unlikely that such tight control of the parameter values would be seen experimentally or physiologically.

(v) *Type V behaviour: excitable behaviour*

For some parameter values, although there is only a single stable equilibrium, the system may exhibit excitability, i.e. relatively small perturbations from the equilibrium can result in a large excursion in phase space before returning and converging back to the equilibrium. In practical terms this means that a large transient spike in the levels of TNF-α and inhibitor can result from an apparently negligible perturbation to the system.

Because the steady state is stable, sufficiently small perturbations from it will always result in the TNF-α and inhibitor concentrations remaining close to steady state concentrations as they eventually settle back to equilibrium. Slightly larger perturbations may, however, result in the dynamics taking the trajectory on a large excursion in phase space before returning to the steady state. We thus see that whether or not we observe a large spike in TNF-α and inhibitor concentrations depends on the precise details of the perturbation to the system. This is illustrated in figure 3*e*, where two initial conditions, close to each other and both close to the equilibrium, result in dramatically different transient behaviours, before both eventually settling down to the equilibrium.

Any real cytokine network will be subject to noise or other apparently random perturbations. One can envisage that in some situations these will occasionally be sufficiently large to result in spikes while at other times they will lead only to small fluctuations about the steady state. The net effect will be to generate an intermittent sequence of spikes at random intervals. This may be another possible explanation for the observed TNF-α spikes during corneal allograft rejection.

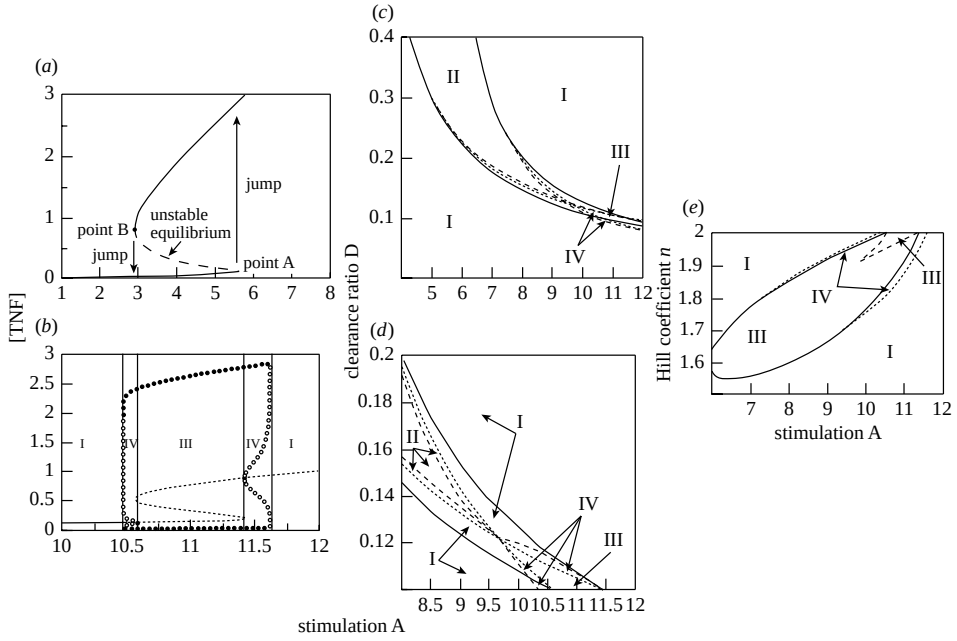

Figure 4. One- and two-parameter bifurcation diagrams determined numerically, showing the ranges of parameter values where type I–IV behaviours exist. (*a*) Threshold behaviour and hysteresis where $B = 0$, $C = 1$, $D = 1$, $E_1 = 0.1$, $E_2 = 0.01$, $F = 1$, $n = 2$; (*b*) zones of oscillatory behaviour, variables as in (*a*); (*c*) two-dimensional bifurcation diagram in *A–D* space, variables as in (*a*); (*d*) a blow-up of (*c*); (*e*) two-dimensional bifurcation diagram in *A–n* space where $B = 0$, $C = 1$, $D = 0.1$, $E_1 = 0.1$, $E_2 = 0.01$, $F = 1$. See text for detailed explanation.

(c) *One-parameter bifurcation diagrams*

The most important parameter in the model is obviously the degree of stimulation, and plotting the concentration of TNF-α against the stimulus allows us to characterize the qualitative changes that may occur in the phase portraits. In particular, how the various behaviours listed above change with varying degrees of stimulation will be explored. A graph showing the equilibrium concentration of TNF-α for a range of values of the stimulation parameter A is known as a bifurcation diagram for the parameter A.

As the stimulation is increased, the equilibrium TNF-α concentration can either change smoothly or become unstable. There are two ways it can destabilize—it can 'jump' to a new and higher steady state concentration at a certain threshold stimulation or it can start to oscillate. Which route of destabilization occurs depends on the other fixed parameters, as will be discussed in § 4(d).

(i) *Equilibrium remains stable*

If the dose response curve for the positive feedback effect of TNF-α on itself is not steep enough (the Hill coefficient is low—see § 4(b)), then there is only a single stable equilibrium for all values of the antigenic stimulation. As the stimulation is increased, the TNF-α concentration will also increase, but nothing unexpected happens.

(ii) *Threshold behaviour*

At low amounts of stimulation, the equilibrium concentration of TNF-α is low, and increases gradually with increasing stimulation (figure 4*a*). However, as the stimulation passes a certain threshold (point A), the equilibrium concentration of TNF-α suddenly 'jumps' to a new and higher level. Technically, this is known as a fold or saddle-node bifurcation. In biological terms, this means that around the threshold point, a small increase in the activation stimulus will result in a sudden large increase in TNF-α. Such behaviour would allow a rapid and strong response to a pathogen once it crosses some antigenic threshold, which makes good evolutionary and physiological sense.

When the stimulation is decreased, there is a different threshold (point B) at which the new equilibrium 'jumps' back to a low steady state concentration of TNF-α. Thus the cytokine network exhibits hysteresis. Hysteresis represents the history dependence of physical systems, and refers to the fact that the system does not return completely to its original state when the stimulation is decreased, showing instead a lag in the values of equilibrium TNF-α concentrations. Reversible hysteresis loops would occur if the stimulation to the system is repeatedly increased and decreased.

This means that the threshold stimulation needed to switch from a high to a low TNF-α equilibrium is lower than that needed to switch from low to a high TNF-α. Put simply, once high levels of TNF-α are induced, it is harder to turn the system off—the activation must be reduced to well below that needed to turn the system on. There will be levels of stimulation at which the system could be in either a high or a low TNF-α level, depending on the previous stimulation history. This would allow the immune system to turn off the response to a pathogen at a lower threshold than that needed for activation, thus ensuring adequate clearance.

Therefore the positive feedback of TNF-α on itself allows the cytokine network to function as a switch between low and high activity levels. The negative feedback loops serve to stabilize its behaviour.

(iii) *Oscillatory behaviour*

In this route (figure 4*b*), the TNF-α concentrations begin to oscillate when the stimulation is increased. Interestingly, the oscillations can begin with small amplitude and grow larger as the stimulation is increased further, or they can emerge suddenly 'fully grown'. The main factor determining which type of behaviour occurs is the slope of the positive feedback dose response, with 'fully grown' oscillations occurring when the slope is steeper.

The first type of oscillation begins when a sink becomes unstable, and a stable cycle is created at a supercritical Hopf bifurcation as the stimulation is increased.

In a supercritical Hopf bifurcation, the amplitude of the created stable cycle starts small and grows larger with increasing stimulation. In the second type of oscillation, a large amplitude stable cycle is initially created at a saddle-node bifurcation of cycles. However, the current equilibrium is still stable with a decreasing basin of attraction as the stimulation increases. The basin of attraction of the sink is bounded by an ever-shrinking unstable cycle, which finally collides with the sink at a subcritical Hopf bifurcation. Trajectories starting outside this boundary are attracted to a large amplitude stable cycle. At the subcritical Hopf bifurcation, the current equilibrium becomes unstable, leaving only the large amplitude cycle created earlier, resulting in the sudden creation of large amplitude oscillations.

Surprisingly, in both cases, as the stimulation is increased further, the oscillations disappear again. Therefore the oscillations only exist for a limited range of antigen concentration, which may explain why such TNF-α oscillations are not commonly observed. We believe that these oscillations are 'side-effects' of the need to ramp up TNF-α production rapidly in response to pathogen, and are unlikely to have a specific function, unlike the case of intracellular calcium oscillations, for example.

(b) *Two-parameter bifurcation diagrams*

Similar to the concept of phase space, parameter space is the multidimensional space with the various parameter variables as axes. At nearly all points in parameter space, the system will have one of the five types of dynamical behaviour listed above. Details of the analysis and calculation of the various bifurcations where one type of behaviour changes into another can be found in standard nonlinear texts, for example Glendinning (1994) or Kuznetsov (1995).

By taking a two-dimensional slice of parameter space, we can study the bifurcation behaviour of the cytokine network with respect to any two parameters of interest. In particular, it shows how stable a particular behaviour (e.g. oscillations) is with respect to change of any two parameters at a time. The two most interesting parameters in addition to the degree of stimulation A are the Hill coefficient n, which represents the steepness of the dose–response curve for positive feedback, and the ratio of clearance of inhibitor to TNF-α, D.

Figure 4*c*,*d* shows the effect of a range of values for A and D on the system behaviour. The most interesting prediction made is that oscillations in such a system cannot arise unless the clearance of TNF-α is relatively faster than that of inhibitor. Such a difference may arise from different decay times or diffusion rates.

Figure 4*e* shows how the dynamic behaviour varies with different values of A and n. Regardless of the degree of stimulation, it shows that the Hill function must have a minimum steepness ($n \approx 1.55$) for oscillations to occur. Below this value of n, there is only a single equilibrium for all values of the stimulation parameter A. Similarly, threshold events also require a minimum steepness for the positive feedback curve, and we indicate regions where one sink or two sinks exist.

If the necessary biological parameters can be experimentally established with sufficient accuracy, then we can use such diagrams to predict what sort of dynamical behaviour to expect if we were to experimentally change various parameters. This has possible use in calculating therapeutic doses of cytokines or cytokine antagonists.

5. DISCUSSION

The role of cytokine networks in the regulation of immune responses is well accepted. This has led to considerable effort in developing new strategies for the treatment of autoimmune diseases, septic shock and the prevention of graft rejection, that rely on disrupting components of the network by use of recombinant cytokines, soluble receptors or antibody (Dallman 1993; Nickerson *et al*. 1997; Kalden *et al*. 1998).

However, the complexity of the network can make it difficult to predict the effect of therapy with such agents. For example, while strategies aimed at blocking TNF-α activity have been successful in the treatment of rheumatoid arthritis and other diseases, they can increase morbidity when used in septic shock (Fisher *et al*. 1996), due to the alterations in the half-life of the cytokine. Models that increase our understanding of the networks may allow both design of the nature, dosage and timing of such therapeutic interventions.

In this paper, we explore a simple cytokine network with just three components (TNF-α, an inhibitor of TNF-α, and an extrinsic stimulation). We demonstrate that even such a simple system can lead to surprisingly rich and even counter-intuitive behaviour, exhibiting threshold effects, hysteresis and oscillating behaviour. Similar pathways are found in many cytokine networks, and therefore the behaviour may be quite widespread.

The original observation that provoked this study was the spikes of TNF-α activity that we observed in a corneal allograft setting. As shown in the model, oscillating behaviour of TNF-α levels can be predicted under certain parameters. We believe that such oscillations in cytokine levels are not directly adaptive, but rather reflect inherent properties of a system designed to produce a rapid response to pathogens. Oscillations in cytokine levels are not unusual in the clinical setting, for example in a spiking fever that may result from changes in TNF-α or other endogenous pyrogens. In addition, juvenile rheumatoid arthritis (Rooney *et al*. 1996) and familial Mediterranean fever (Schattner *et al*. 1991) are associated with spikes in temperature and cytokine levels. An understanding of any positive and negative feedback pathways involved in such oscillations may allow better treatment of these diseases.

The model makes several specific and surprising predictions about the behaviour of such a cytokine network under different degrees of antigenic stimulation which may be tested experimentally or observed in clinical settings.

One such prediction is that a reduction in the causal stimulus may push the system from having a stable equilibrium concentration of TNF-α to having oscillatory behaviour, where both the peak and mean levels of cytokine may be higher than pretreatment. Clinically, this might mean that in some cases inadequate treatment may be worse than no treatment at all.

Another prediction is that with increasing antigenic or mitogenic stimulation to the cytokine network, we should

see threshold and hysteresis behaviour. Clearly this could be of adaptive benefit, resulting in a threshold of stimulation that causes a switch from low to high TNF-α production. This again has a number of clinical parallels; small changes in an underlying infection can lead to sudden changes in the condition of a patient. At points close to this threshold, small alterations can have a dramatic effect, which is clearly important when devising new treatment strategies.

The hysteresis in the response curve ensures that it is difficult to turn off a response once it has started—in other words, a stimulus below the threshold necessary to lead to high equilibrium levels, will be capable of maintaining the system in the high equilibrium state. These threshold and hysteresis effects are recognized features of many inflammatory diseases; small changes in either the stimulus or some of the parameters may result in a threshold change leading to an inflammatory state. When the threshold has been crossed the system is 'locked in' the inflammatory state and requires large changes in the conditions such as high doses of anti-inflammatory drugs before it can revert to the quiescent state. This may also be related to 'flare ups' seen in many autoimmune diseases.

A third prediction is that for oscillations to arise in such a cytokine network, the slope of the positive feedback dose–response curve must exceed a minimal steepness. In addition, the clearance rate TNF-α must be faster than the clearance rate of inhibitor. These two conditions can be seen from figure 4*e*,*c*, respectively.

While the model is undoubtedly a simplification of the actual process, the assumptions made are biologically justifiable, and its predictions should be experimentally verifiable. The weakness of the model is that reliable estimates of the relevant parameters are not currently available, though there appears no technical reason why they cannot be experimentally determined. Even with purely qualitative predictions, the model already serves the useful function of suggesting new phenomena to look out for, and possible new experiments to bring out such behaviour.

When the kinetic data are available, it will be possible to make quantitative predictions, and also to modify or extend the model in the light of experimental findings. Eventually, it is possible to envisage a 'virtual laboratory' in which one can test the effect of various interventions and their combinations on cytokine networks *in silico*, which would be prohibitively time-consuming and expensive in the immunological laboratory. The promising interventions identified in this way can then be tested conventionally, and hopefully result in useful therapeutics.

REFERENCES

Amiot, F., Boussadia, O., Cases, S., Fitting, C., Lebastard, M., Cavaillon, J.-M., Milon, G. & Dautry, F. 1997 Mice heterozygous for a deletion of the tumor necrosis factor-α and lymphotoxin-α genes: biological importance of a nonlinear response of tumor necrosis factor-α to gene dosage. *Eur. J. Immunol.* **27**, 1035–1042.

Dallman, M. J. 1993 Cytokines as mediators of organ graft rejection and tolerance. *Curr. Opin. Immunol.* **5**, 788–793.

Ferguson, T. A., Fletcher, S., Herndon, J. & Griffith, T. S. 1995 Neuropeptides modulate immune deviation induced via the anterior chamber of the eye. *J. Immunol.* **155**, 1746–1756.

Ferguson, T. A., Herndon, J. M. & Dube, P. 1994 The immune response and the eye: a role for TNFα in anterior chamber-associated immune deviation. *Invest. Ophthalmol. Vis. Sci.* **35**, 2643–2651.

Fisher, C. J., Agosti, J. M., Opal, S. M., Lowry, S. F., Balk, R. A., Sadoff, J. C., Abraham, E., Schein, R. M. & Benjamin, E. for the Soluble TNF Receptor Sepsis Study Group 1996 Treatment of septic shock with the tumor necrosis factor receptor:Fc fusion protein. *New Engl. J. Med.* **334**, 1697–1702.

Glendinning, P. 1994 *Stability, instability and chaos: an introduction to the theory of nonlinear differential equations*, Cambridge texts in applied mathematics. Cambridge University Press.

Griffith, T. S., Brunner, T., Fletcher, S. M., Green, D. R. & Ferguson, T. A. 1995 Fas ligand-induced apoptosis as a mechanism of immune privilege. *Science* **270**, 1189–1192.

Heaney, M. L. & Golde, D. W. 1996 Soluble cytokine receptors. *Blood* **87**, 847–857.

Kalden, J., Breedveld, F., Burkhardt, H. & Burmester, G. 1998 Immunological treatment of autoimmune diseases. *Adv. Immunol.* **68**, 333–418.

Kuznetsov, Y. A. 1995 *Elements of applied bifurcation theory*, applied mathematical sciences series. Berlin: Springer.

Nickerson, P., Steiger, J., Zheng, X., Steele, A., Steurer, W., Roy-Chaudhury, P. & Strom, T. 1997 Manipulation of cytokine networks in transplantation: false hope or realistic opportunity for tolerance? *Transplantation* **63**, 489–494.

Niederkorn, J. Y. 1990 Immune privilege and immune regulation in the eye. *Adv. Immunol.* **48**, 191–226.

Philip, R. & Epstein, L. B. 1986 Tumour necrosis factor as immunomodulator and mediator of monocyte cytotoxicity induced by itself, γ-interferon and interleukin-1. *Nature* **323**, 86–89.

Prieur, A., Lombard, P. R. & Dayer, J. 1996 Dynamics of fever and the cytokine network in systemic juvenile arthritis. *Rev. Rhum. Engl. Ed.* **63**, 163–170.

Rayner, S., King, W., Isaacs, J., Hale, G., George, A. & Larkin, D. 1999 Pulsatile local bioactive tumor necrosis factor in allograft rejection. *Invest. Ophthalmol. Vis. Sci.* (Submitted.)

Rooney, M., David, J., Symons, J., Giovine, F. D., Varsani, H. & Woo, P. 1995 Inflammatory cytokine responses in juvenile chronic arthritis. *Br. J. Rheumatol.* **34**, 454–460.

Schattner, A., Lachmi, M., Livneh, A., Pras, M. & Hahn, T. 1991 Tumour necrosis factor in familial mediterranean fever. *Am. J. Med.* **90**, 434–438.

Segel, L. A. 1984 *Modeling dynamic phenomena in molecular and cellular biology.* Cambridge University Press.

Streilein, J. W., Wilbanks, G. A. & Cousins, S. W. 1992 Immunoregulatory mechanisms of the eye. *J. Neuroimmunol.* **39**, 185–200.

Strieter, R. M., Kunkel, S. L. & Bone, R. C. 1993 Role of tumor necrosis factor-α in disease states and inflammation. *Crit. Care Med.* **21**, S447–S463.

Tracey, K. J. & Cerami, A. 1994 Tumour necrosis factor: a pleiotropic cytokine and therapeutic target. *A. Rev. Med.* **45**, 491–503.

Williams, K. A. & Coster, D. J. 1997 Rethinking immunological privilege: implications for corneal and limbal stem cell transplantation. *Mol. Med. Today* **11**, 495–501.

Williams, K., Roder, D., Esterman, A., Muehlberg, S. & Coster, D. 1992 Factors predictive of corneal graft survival. report from the Australian Corneal Graft Registry. *Opthalmology* **99**, 403–414.

As this paper exceeds the maximum length normally permitted, the authors have agreed to contribute to production costs.