M3. Molecular resolution module

M3.1 Slot, options and selection

What the slot controls

Slot M3 selects the resolution at which four signalling systems are represented: TNF with its receptors TNFR1 and TNFR2; IL-10 with its receptor IL-10R; the intracellular NF-κB pathway; and MHC class II antigen presentation.

Coarse-grained option (default). Receptor occupancy is computed in closed form from the local field value and the agent's receptor count, and receptor-mediated consumption is an algebraic sink. No per-agent ordinary differential equation is solved. The option is specified in S6 (S6.2 to S6.13).

Fine-grained option. Each of the four systems can instead be represented by a block of ordinary differential equations integrated per agent on the molecular substep. The four components are selected independently: TNF receptor dynamics (M3.3), IL-10 receptor dynamics (M3.4), NF-κB dynamics (M3.5) and MHC class II dynamics (M3.6). Partial selections coexist with the coarse-grained relations for the unselected systems.

All four components are off by default, so no per-agent ODE is solved at any point. This is the basis for the spine's statement that no nested ODD is required (spine 1.1). The molecular substep is retained because the coarse-grained sinks act at that cadence (S6.1).

Selection

Four boolean simulation flags select the components, one per system, each set from its own command-line option (Appendix B). All four are initialized false in the simulation constructor. They are serialized with a saved state.

The flags are passed to agents as call arguments, not held per agent. There is therefore no per-agent resolution state, and all agents of one type run at one resolution. Resolution can differ between cell types (M3.2).

A study selecting the fine-grained option would need to sample a different parameter set. The sampled coarse-grained thresholds and rate constants (`Core_thresholdApoptosisTNF`, `Core_kApoptosis`, `Core_Mac_thresholdNFkBTNF`, `Core_Mac_kNFkB`, `Core_Mac_thresholdSTAT3IL10`, `Core_Mac_kSTAT3IL10`) govern macrophage decisions only under the coarse-grained option.

A macrophage carries the full complement of molecular state variables whatever the selection. These variables are initialized and serialized in every configuration; they are used only by the enabled components. Each species' surface receptor counts remain inputs to the coarse-grained relations whenever that species' receptor component is off.

Schedule

Block	When	Specified in
Parameter resolution	Once at startup, after the diffusion method is set and before initialization	M3.6 (MHC class II rescaling and derivations)
Agent construction	Whenever an agent is created by a constructor	S6.3 (receptor draws); M3.5 (NF-κB equilibration)
Secretion	In each split half-step, twice per diffusion step	M3.3, M3.4, M3.5
Molecular substep loop	N_m/2 substeps in each split half-step	M3.2
Next-state computation	Once per cellular timestep, per agent	M3.6 (T cell activation), M3.7 (decision points)

The two gates in the substep loop

The substep loop makes two calls. They are not mutually exclusive.

```
solve <=> tnfr or il10r or nfkb or mhc  integrate the per-agent ODEs (M3.2)
sink  <=> !il10r or !tnfr               apply the coarse-grained sinks (S6.8)
```

Within the sink routine each species' sink is re-tested on its own flag. The effective rule is that the TNF sink is applied if and only if the TNF receptor flag is clear, and the IL-10 sink if and only if the IL-10 receptor flag is clear. The outer condition only skips the call when both inner tests would fail. With exactly one receptor component selected, the integrator and one sink run in the same substep. Neither sink is subject to the M8 depletion flags, so depletion suppresses production, not consumption.

The rule has one exception, created by the NF-κB failsafe (M3.2): with NF-κB dynamics and without TNF receptor dynamics, TNF is depleted by both mechanisms (M3.8).

Distinct combinations

The macrophage dimension excludes the anti-TNF blocks of M7. T cell dimensions are given in M3.2.

Components selected	Integrator runs	TNF sink	IL-10 sink	Macrophage dimension	Note
None	No	Yes	Yes	0	Default
TNFR	Yes	No	Yes	10	

Components selected	Integrator runs	TNF sink	IL-10 sink	Macrophage dimension	Note
IL-10R	Yes	Yes	No	3	
TNFR, IL-10R	Yes	No	No	13	The outer condition short-circuits
MHC II only	Yes	Yes	Yes	16	Both coarse-grained sinks retained
NF-κB	Yes	Yes	Yes	36	TNFR block forced; TNF consumed by block and sink (M3.8)
NF-κB, IL-10R	Yes	Yes	No	39	TNFR block forced; TNF consumed by block and sink (M3.8)
NF-κB, TNFR	Yes	No	Yes	36	
NF-κB, TNFR, IL-10R	Yes	No	No	39	
All four	Yes	No	No	55	

Adding the MHC class II component to any row adds 16 to the macrophage dimension and leaves the sink columns unchanged, since that block neither reads nor writes the TNF or IL-10 field.

What each fine-grained component replaces in S6

Component	S6 statements that cease to apply
TNF receptors	S6.4, S6.5, S6.6; the TNF term of S6.8; the field write of S6.11
NF-κB pathway	Those of TNF receptors, and the chemokine and TNF parts of S6.9 to S6.13
IL-10 receptors	S6.7; the IL-10 term of S6.8; the field write of S6.12
MHC class II	None

The unit conventions and conversions (S6.2, "Field units"), the degradation constants (S6.15) and the derived-field relation (S6.14) hold under every selection, the last subject to the literal 1 : 1 : 2 of the NF-κB chemokine output (M3.5).

Several S6 results rest on arguments that do not survive a fine-grained component and must be re-derived under it:

- occupancy as an equilibrium quantity with no memory and no receptor depletion (S6.2), since fine-grained receptor species are integrated and depleted by binding;
- the ceilings of the signalling quantity and the apoptosis test (S6.4, S6.6), since the fine-grained branches clamp different quantities against different denominators and their own parameters;
- linearity of the TNF tests in the occupancy (S6.5, S6.6);
- the forms of the two IL-10 secretion inhibitions (S6.10, S6.11), since under NF-κB dynamics neither helper is called for macrophages;
- secretion per timestep as rate $\times$ state factor $\times$ interval (S6.9), since secretion becomes the output of an integration;
- the sink scope (S6.8), since each sink applies only for the species whose receptor component is off;
- the derived-field relation's requirement that every chemokine source writes the three chemokines in the proportion of the rate ratios (S6.14), which holds under NF-κB dynamics because the literal 1 : 1 : 2 equals the default rate ratios.

Rate variables S6 sets for M3

The secretion helpers set per-agent rate variables in every configuration. A fine-grained block consumes them; otherwise they are stored and never read.

Variables	Set by	Consumed by
TNF transcription ceiling and floor	TNF secretion helper	TNF receptor block (M3.3)
IL-10 synthesis rate	IL-10 secretion helper	IL-10 receptor block (M3.4)
Three NF-κB transcription coefficients	NF-κB secretion helper	NF-κB block (M3.5)

The state factors of S6.13 therefore remain the specification of relative secretory capacity under both resolutions. The exception is the macrophage's chemokine and TNF output under NF-κB dynamics, where the transcription arguments of M3.5 take their place.

Reads and writes

Reads:

- TNF, IL-10 and shed TNFR2 at the agent's own compartment (M3.3, M3.4);

- the M7 drug and complex fields at the agent's own compartment (M7);
- per-agent receptor, bound, internalized, transcript and protein species;
- the per-agent rate variables set by secretion (read only);
- the MHC class II neighbourhood inputs (M3.6);
- the agent's own intracellular bacterial count (M3.6);
- the M8 TNF and IL-10 depletion flags.

Writes:

- TNF, IL-10 and shed TNFR2 at the agent's own compartment, by assignment (M3.3, M3.4);
- the M7 drug and complex fields, by assignment (M7);
- the per-agent receptor, bound, internalized, transcript and protein species, in place;
- CCL2, CCL5 and CXCL9 at the agent's own compartment, by increment (M3.5);
- T cell state, from the macrophage next-state pass (M3.6);
- the observation quantities, in particular surface antigen-loaded MHC class II (M3.6).

M3.2 Per-agent ODE machinery

This section holds the machinery shared by all four components: which agents are integrated, how the equation system is laid out, how it is built, stepped and exchanged with the grid. M7.8 refers to this section for the same machinery.

Which agents are integrated

The substep loop iterates fixed lists of agent types, not every agent.

- **Macrophages** are integrated whenever any of the four components is selected.
- **T cells** (Tgam, Tcyt, Treg) are integrated when TNF receptor, IL-10 receptor or NF-κB dynamics is selected. Under MHC class II dynamics alone the T cell loop is not entered.
- **Dendritic cells and neutrophils** appear in neither the integration loop nor the sink loop. The exclusion is by an enumerated list, not by a guard. Neutrophils are specified in M9 (M9.7); the sink part of the exclusion is recorded at S6.1.

The integration loop is guarded on the simulation flags, not on the resulting equation dimension, so an agent whose layout has dimension zero is still iterated.

Block layout

The derivative function records, for each of six possible blocks, the offset at which the block begins in the agent's solution vector. Blocks are admitted in a fixed order, contiguous and without padding; an absent block is absent from the vector, and the dimension is the sum of the admitted sizes.

Block	Equations	Admitted when
TNFR	10	TNF receptor dynamics, or NF-κB dynamics by the failsafe
IL-10R	3	IL-10 receptor dynamics
NF-κB	26	NF-κB dynamics
MHC class II	16	MHC class II dynamics
Etanercept	3	M7, Etanercept dosed
Infliximab	7	M7, Infliximab dosed

Two cross-block couplings exist. The anti-TNF equations modify the membrane TNF equation of the TNFR block (M7). TNFR transcription reads bound IL-10R from the IL-10R block when that block is present (M3.3).

The NF-κB failsafe

The derivative function's constructor sets its TNFR argument true whenever its NF-κB argument is true, because the NF-κB equations read bound surface TNFR1 and write membrane TNF, which exist only in the TNFR block.

The assignment is to the constructor's local copy. The simulation's TNF receptor flag stays false everywhere else in the program. The consequences are collected in M3.8.

Dispatch

```
if NF-kB block present:     combined NF-kB and TNFR routine (writes both blocks; M3.5)
else if TNFR block present: TNFR-only routine (also writes the anti-TNF blocks; M3.3, M7)
if IL-10R block present:    IL-10R routine (M3.4)
if MHC block present:       MHC class II routine (M3.6)
```

The NF-κB and TNFR-only branches are exclusive. Under NF-κB dynamics the TNFR equations are written by different code and differ (M3.5). The IL-10R and MHC class II routines are independent of them and of each other.

Derivative function and stepper

The derivative function is obtained through a per-class accessor. The base agent class's accessor forces the NF-κB and MHC class II flags false; only the macrophage class overrides it, passing all four flags through. T cells therefore use the base class's function, which never carries the NF-κB or MHC class II block.

Each class builds its derivative function once, on that class's first solve, and caches it for the rest of the process. Its layout is fixed by the flags in force at that first call. The macrophage class holds its own function and stepper; the three T cell types share the base class's pair. This is how two layouts coexist in one simulation. A flag changed mid-run would leave the old layout integrated while the decision points of M3.7 select branches for the new flags; batch runs never change flags.

Agents of a class share one stepper. The stepper is rebuilt only if the requested integration method differs from that at its first construction. A derivative entry that the current evaluation does not write keeps the value left in the shared buffer by the most recent evaluation in that class that wrote it, and is zero only if no evaluation in that class writes it. This bears on the two imposed MHC class II positions (M3.6) and on the anti-TNF entries under NF-κB dynamics (M7.8).

The integration method is selected by the `--odesolver` command-line option, whose program default is method 4, fourth-order Runge–Kutta. After construction the method holds the enumeration's count sentinel; the setter then runs from the option on every run and every load, so the sentinel never reaches a timestep.

Substep subdivision under NF-κB dynamics

When NF-κB dynamics is selected, a capable agent advances each molecular substep in `NFkB_NFkBTimeCoefficient` steps of Δt_m / `NFkB_NFkBTimeCoefficient`. Total time per substep is unchanged.

Capability is a per-class predicate. The macrophage class reports itself capable; the base class reports itself incapable, and the T cell classes do not override it. Subdivision therefore applies to macrophages only. Marshalling ("Exchange with the grid") occurs once per molecular substep, outside the subdivision loop.

Cell-type scope

Cell type	TNFR	IL-10R	NF-κB	MHC II	Sink applied	Subdivision
Macrophage	Yes	Yes	Yes	Yes	Yes	Yes
Tgam, Tcyt, Treg	Yes	Yes	No	No	Yes	No
Dendritic cell	No (not iterated)	No	No	No	No	n/a
Neutrophil	No (not iterated)	No	No	No	No	n/a
Fibroblast, myofibroblast	Separate routine (M4)				No	n/a
Mycobacterium	No	No	No	No	No	n/a

The dendritic cell class derives from the macrophage class and inherits its accessor and capability predicate, so it would run all four components if iterated. It is not iterated. Both dendritic cells and neutrophils nonetheless take the receptor and NF-κB construction draws (S6.3), so they carry the state without the dynamics. Neither type is created by default, so their exclusion has no effect in the default configuration; it is live code, not dead code.

T cells have no NF-κB state variable and no NF-κB-gated transition, run no construction-time equilibration (M3.5), and evaluate TNF-induced apoptosis with the NF-κB argument fixed false (M3.7).

Components selected	Macrophage dimension	T cell dimension
TNFR	10	10
IL-10R	3	3
TNFR, IL-10R	13	13
MHC II	16	Not integrated
NF-κB	36	0
NF-κB, IL-10R	39	3
All four	55	13

Under NF-κB dynamics alone, T cells receive a function with no block: the base accessor forces the NF-κB flag false and the failsafe then acts on the forced value. The T cell loop still executes, on a vector of length zero. A loop guarded on the dimension would skip the call instead, which is not identical behaviour.

Exchange with the grid

Before each integration the agent's solution vector is filled from its own members and from the grid at its own compartment; after it, the results are copied back to the same places. The within-block mapping is fixed and is the same in both directions.

Nine positions, across the TNFR, IL-10R and anti-TNF blocks, hold field values at the agent's own compartment, not per-agent state. They are read before the substep and written back after it by assignment. For such a species the integrator owns the field value at that compartment for the substep. The implementation observes these rules:

- all production, binding and internalization of that species there must be in the block;
- the coarse-grained secretion increment for that species is suppressed (M3.3, M3.4);
- the coarse-grained sink for that species is suppressed (M3.1).

No other process writes an assigned field within the substep: secretion and the diffusion solve act outside the substep loop. The one exception is the coarse-grained TNF sink under NF-κB dynamics without TNF receptor dynamics (M3.8).

Agents sharing a compartment are processed in list order. Each reads the value its predecessor wrote back, so the field is integrated sequentially, not simultaneously (spine Element 3). A reimplementation must not aggregate the writes of co-located agents. It must also apply no separate degradation step to an assigned field within the substep; decay is folded into the diffusion solve (M1).

A field written by increment carries no such constraint (M3.5).

Failure behaviour

On copy-back, a surface antigen-loaded MHC class II count that is not a number throws an exception identifying the agent and its position, and the run ends. Under the imposed-input formulation of M3.6 this check can end a run (M3.6, Rescaling at load time). No other quantity is checked or clamped, so negative concentrations from a long substep are carried forward.

M3.3 TNF receptor dynamics

Summary. TNF binding to TNFR1 and TNFR2, internalization, recycling, shedding and membrane TNF release are integrated per agent in a ten-equation block, which owns the TNF and shed TNFR2 field values at the agent's compartment for each substep. Receptor pools relax towards each agent's drawn counts, and IL-10 modulates TNF transcription through a logistic in bound IL-10R. The TNF secretion increment and the coarse-grained TNF sink are suppressed.

Notation

M_r is the TNF transcript and M membrane TNF; R_1, R_2 free surface TNFR1 and TNFR2; C_1, C_2 the bound surface complexes; I_1, I_2 the internalized bound complexes. These are in molecules per cell. T is soluble TNF and S the shed TNFR2 complex, both in molecules per compartment (spine 2.1).

Every second-order association term uses the molar concentration, and the dissociation rate constants are formed from the association and equilibrium constants:

$$T_M = \frac{T}{N_A V}, \qquad k_{\mathrm{off},1} = k_{\mathrm{on},1} K_{D,1}, \qquad k_{\mathrm{off},2} = k_{\mathrm{on},2} K_{D,2}.$$

The ten equations

$$\frac{dM_r}{dt} = k_{\mathrm{mRNA}}^{\mathrm{eff}} - k_{\mathrm{trans}} M_r$$

$$\frac{dM}{dt} = k_{\mathrm{trans}} M_r - k_{\mathrm{TACE}} M$$

$$\frac{dR_1}{dt} = v_1 - k_{\mathrm{on},1} T_M R_1 + k_{\mathrm{off},1} C_1 - k_{T,1} R_1 + k_{\mathrm{rec},1} I_1$$

$$\frac{dR_2}{dt} = v_2 - k_{\mathrm{on},2} T_M R_2 + k_{\mathrm{off},2} C_2 - k_{T,2} R_2 + k_{\mathrm{rec},2} I_2$$

$$\frac{dC_1}{dt} = k_{\mathrm{on},1} T_M R_1 - k_{\mathrm{off},1} C_1 - k_{\mathrm{int},1} C_1$$

$$\frac{dC_2}{dt} = k_{\mathrm{on},2} T_M R_2 - k_{\mathrm{off},2} C_2 - k_{\mathrm{int},2} C_2 - k_{\mathrm{shed}} C_2$$

$$\frac{dI_1}{dt} = k_{\mathrm{int},1} C_1 - k_{\mathrm{deg},1} I_1 - k_{\mathrm{rec},1} I_1$$

$$\frac{dI_2}{dt} = k_{\mathrm{int},2} C_2 - k_{\mathrm{deg},2} I_2 - k_{\mathrm{rec},2} I_2$$

$$\frac{dT}{dt} = \phi \left(k_{\mathrm{TACE}} M - k_{\mathrm{on},1} T_M R_1 + k_{\mathrm{off},1} C_1 - k_{\mathrm{on},2} T_M R_2 + k_{\mathrm{off},2} C_2 \right)$$

$$\frac{dS}{dt} = \phi \, k_{\mathrm{shed}} C_2$$

The factor $\phi = (\rho_{\mathrm{cell}}/N_A) \cdot N_A V = \rho_{\mathrm{cell}} V$, with $\rho_{\mathrm{cell}} = 1.25 \times 10^{11}$ cells per litre and $V = 8 \times 10^{-12}$ L, so $\phi = 1$: one cell per compartment. The field equations are therefore direct molecule-per-compartment bookkeeping, and ϕ may be omitted. It is unity only through the values of two constants, so changing either alone rescales all receptor-mediated TNF exchange.

Conservation and consumption

The binding terms of the soluble TNF equation are the exact negatives of those of the two bound-complex equations, so binding conserves TNF. TNF is consumed only by internalization. The internalized complex is degraded or recycled, and recycling returns the receptor to the free pool but returns no TNF. Internalization is what the coarse-grained TNF sink substitutes for: `Core_estConsRateTNF` equals `Molecular_TNF_kInt1`, both 7.7×10^{-4} s^{-1} by default.

Shed TNFR2

Shedding moves bound TNFR2 into its own field, and no term in the block returns anything from it. The shed field is read only as the block's tenth position, so the value written back exceeds the value read by the amount shed. It has its own diffusivity, `Molecular_TNF_diffusivityShedTNFR2`, and is advanced by the diffusion module. The diffusion step, gated on the TNF receptor flag, dissociates the shed complex and returns TNF to the field (M1.3, "Post-solve binding sweep").

In that step the complex is decremented first, and the TNF returned is computed from the already-decremented value, so slightly less TNF returns than the complex loses (M1.3, "Post-solve binding sweep").

The shed field is allocated, maintained and serialized in every configuration, and is identically zero by default.

Receptor synthesis and the set-point

The receptor synthesis rates are fixed per agent at construction as v_1 = R_1(0) k*{T,1} and v_2 = R_2(0) k*{T,2}, from the counts drawn at construction (S6.3). Without TNF the free-receptor equation reads $dR_1/dt = k_{T,1}(R_1(0) - R_1)$, so each pool relaxes to the agent's own drawn value. The drawn count is a set-point, not only an initial condition: heterogeneity survives, and a depleted cell recovers towards its own value. The synthesis rates are never recomputed and ignore signalling state.

Under this block the free surface counts that S6 reads as fixed agent properties become state variables that change over the cell's life.

Transcription and IL-10 modulation

Let B be the bound IL-10R count. It is read from the IL-10R block when that block is present. Otherwise it is estimated from the local IL-10 field,

$$B = \frac{c_{\text{IL-10}}\,\bar{R}_{\text{IL-10}}}{K_{D,\text{IL-10}} + c_{\text{IL-10}}}, \qquad c_{\text{IL-10}} = \frac{C_{\text{IL-10}}}{N_A V},$$

where $\bar{R}_{\text{IL-10}}$ is the agent's stored type-mean IL-10R count, not its drawn count. So without the IL-10R block, IL-10 sensing in this term is homogeneous within a cell type, while the coarse-grained IL-10 sink, which uses the drawn count, is heterogeneous (S6.8).

The effective transcription rate is

$$k_{\text{mRNA}}^{\text{eff}} = k_{\text{synth}} + \frac{k_{\text{mRNA}} - k_{\text{synth}}}{1 + \exp\left(\dfrac{B - \gamma}{\delta}\right)},$$

with γ = `Molecular_IL10_LinkRNAGamma` and δ = `Molecular_IL10_LinkRNADelta`. It interpolates from k_mRNA at negligible IL-10R occupancy to k_synth at saturation, centred at B = γ with width δ. With $\gamma = 20$ and $\delta = 7$, the transition from nine-tenths to one-tenth of the range runs from 4.6 to 35.4 bound receptors, 0.31 % to 2.36 % of a mean complement of 1500. It is in effect a switch on the absolute number of occupied receptors.

k_mRNA and k_synth are the two per-agent rate variables set by the TNF secretion helper from `Molecular_TNF_Mac_kRNAMac` and `Molecular_TNF_Mac_kSynthMac` (T cells: `Molecular_TNF_Tcell_kRNATcell`, `kSynthTcell`). No term adds membrane TNF in proportion to k_synth: membrane TNF is produced only by translation. k_synth is the IL-10-saturated floor of transcription, and k_mRNA its ceiling. Their ratio is 0.14 for both cell types, equal to `Core_MaxIL10Inhibition`, which no code reads (S6.11).

The expression is evaluated only when k_mRNA > 0; otherwise the rate is zero, which silences non-secreting cells. Under full TNF depletion the rate is zero; under fractional TNF depletion it is multiplied by `Core_depletionFractionTNF` (M8).

Under NF-κB dynamics six of these equations are written differently (M3.5). Under M7 the membrane TNF equation gains complex-formation and dissociation terms, written by this same routine, which is why the anti-TNF blocks need the TNFR block (M7.8).

TNF secretion and the coarse-grained sink

The TNF secretion helper computes a field increment and the two rate variables in every configuration. The increment is applied only when the TNF receptor flag is clear; the rate variables are consumed only by this block. In a coarse-grained configuration they are computed, stored and never read.

- The coarse-grained IL-10 logistic in $\log_{10}$ field (S6.11) multiplies the increment only, so it has no effect under this block. Its fine-grained counterpart is the transcription logistic of M3.3, a different form acting on a different quantity.
- The macrophage state factors depend on the cell's discrete NF-κB flag (S2.4), not on the NF-κB component flag (S6.13).
- The Tcyt TNF state factor is 0.1 and its construction-time synthesis argument a tenth of the Tgam value, so the reduction applies to both products.
- The producer test gates the TNF helper entirely. A cognate non-producer keeps its construction values (k_mRNA = 0; k_synth = `Molecular_TNF_Tcell_kSynthTcell` for Tgam, a tenth of it for Tcyt) until first suppressed, and transcribes no TNF because k_mRNA is zero.
- Under M4 the TNF state factor is multiplied by a factor in cumulative bound TGF-β, floored at one half; it scales the rate variables as well as the increment (M4).

The coarse-grained TNF sink (S6.8) follows the gate and the depletion exemption of M3.1. Both overrides of the sink routine pass the agent's own receptor counts, so the sink is heterogeneous.

Fields written

Field	Mode	Default
TNF	Assignment	Written coarsely (S6)
Shed TNFR2	Assignment	Allocated, identically zero

M3.4 IL-10 receptor dynamics

Summary. IL-10 binding to IL-10R and internalization of the bound receptor are integrated per agent in a three-equation block, which owns the IL-10 field value at the agent's compartment for each substep. IL-10 secretion enters the field equation as a per-agent synthesis rate, and the receptor pool relaxes towards each agent's drawn count. The IL-10 secretion increment and the coarse-grained IL-10 sink are suppressed.

The three equations

L is soluble IL-10 in molecules per compartment; R free and B bound surface IL-10R in molecules per cell. As in M3.3, k_off = k_on K_D, L_M = L / (N_A V) and $\phi = 1$.

$$\frac{dL}{dt} = \phi\,(k_{I,\mathrm{synth}} + k_{\mathrm{off}} B - k_{\mathrm{on}} R\, L_M)$$

$$\frac{dR}{dt} = v_I - k_{\mathrm{on}} R\, L_M + k_{\mathrm{off}} B - k_{I,T} R$$

$$\frac{dB}{dt} = k_{\mathrm{on}} R\, L_M - k_{\mathrm{off}} B - k_{I,\mathrm{int}} B$$

The structure parallels one TNF receptor subtype, with these differences:

- there is no transcript or membrane form, secretion entering the field equation directly as k_{I,synth};
- there is no recycling, so internalization is a pure loss of ligand and receptor, and no internalized species is tracked;
- v_I = R(0) k_{I,T} is fixed at construction, so the pool relaxes to the agent's drawn count, as in M3.3.

Binding conserves IL-10, which is consumed only by internalization.

Synthesis and depletion

k_{I,synth} is the per-agent IL-10 synthesis rate set by the IL-10 secretion helper from the state factor and the per-type rate argument; the block only reads it. All IL-10 depletion is applied upstream, in that helper. A reimplementation should deplete in secretion, not in the derivative.

Reads by other blocks

The integrated bound count B is read by the TNFR transcription logistic when this block is present (M3.3) and by the STAT3 test (M3.7). Both read the count itself, not an occupancy fraction.

IL-10 secretion

The IL-10 secretion helper is called on every macrophage branch, including the NF-κB branch, with the same state factors. Its rate argument by macrophage state:

Macrophage state	IL-10 synthesis argument
Resting	0
Infected, chronically infected	`Molecular_IL10_Mac_IkSynthMacInf`
Activated	`Molecular_IL10_Mac_IkSynthMacAct`, then multiplied by the plasticity factor

The T cell argument is `Molecular_IL10_Tcell_IkSynthTcell`. The IL-10 state factors of Tgam and Tcyt are zero (S6.13), so IL-10 is inert for effector T cells at both resolutions and only the Treg secretes it.

The field increment is applied when the IL-10 receptor flag is clear and no IL-10 depletion is in force. Otherwise it is applied when fractional IL-10 depletion is in force, without re-testing the receptor flag. So with this block selected and fractional IL-10 depletion in force, IL-10 is secreted both by increment and through k_{I,synth}. M8 specifies this path.

The activated macrophage's IL-10 induction by TNF uses the coarse-grained occupancy estimate in both resolutions, whatever the TNF receptor flag (S6.12). It is the one quantity with a fine-grained counterpart that is always computed coarsely.

After the helper returns, the activated macrophage's IL-10 synthesis rate is multiplied by C_1 / (C_1 + `Molecular_IL10_Mac_halfSatIL10mRNA`), with C_1 the bound surface TNFR1 count, in both secretion branches. The rationale is that it represents polarization plasticity. C_1 is zero unless the TNFR block is admitted, so the factor is zero without it. That is inert by default, but under IL-10 receptor dynamics without TNF receptor dynamics activated macrophages synthesize no IL-10, while infected and chronically infected macrophages are unaffected.

The coarse-grained IL-10 sink

The coarse-grained IL-10 sink (S6.8) follows the gate and the depletion exemption of M3.1.

Fields written

Field	Mode	Default
IL-10	Assignment	Written coarsely (S6)

M3.5 NF-κB dynamics

Summary. The intracellular NF-κB pathway is integrated per macrophage in a 26-equation block, driven by bound surface TNFR1 and always accompanied by the TNFR block. Its outputs export chemokine to the grid and TNF to the membrane, and its activation and apoptosis-inhibitor levels replace the coarse-grained quantities in two macrophage decisions (M3.7). Each macrophage equilibrates the pathway for twelve simulated hours at construction.

Relationship to the TNFR block

The NF-κB block cannot be admitted alone; the TNFR block is forced in (M3.2). The combined routine writes both blocks, and its TNFR equations differ from M3.3 in three ways.

The transcript is frozen. Its derivative is zero unconditionally, so the transcript stays at its initial value of zero and the transcription logistic of M3.3 is not evaluated. Membrane TNF is produced by export from the pathway's intracellular TNF:

$$\frac{dM}{dt} = e_{3,\mathrm{TNF}}\,\mathrm{TNF}_{\mathrm{int}} - k_{\mathrm{TACE}} M.$$

A fixed loss is added to dissociation. The two free-receptor and the two bound-complex equations use $k_{\mathrm{off},i} + \kappa$ in place of k*{off,i}, where κ, a compiled constant of* 4.58×10^{-4} *s*$^{-1}$*, is not a parameter. The soluble TNF equation keeps k*{off,i} alone. TNF is therefore not conserved: a complex lost at rate κ returns its receptor to the free pool and its TNF nowhere. κ is 3 % of k*{off,2} but 86 % of $k*{\text{off},1} = 5.32 \times 10^{-4}$ s^{-1}, so about 46 % of the events returning bound TNFR1 to the free pool return no TNF. A reimplementation applying κ symmetrically would not reproduce the model. The internalized-complex, soluble TNF and shed TNFR2 equations are as in M3.3, and the combined routine contains no anti-TNF terms.

IL-10 is not coupled to TNF production. The routine computes the equilibrium bound IL-10R count and the depletion factors but applies neither to any TNF term, because no such linkage is established for this model. IL-10 therefore acts on macrophage TNF production under coarse-grained and TNF-receptor-only resolution, not under NF-κB dynamics.

Receptor-proximal signalling

K_a is active IKK kinase; K_n, K^+, K^- neutral, active and inactive IKK; A is A20. $\mathcal{K}$ = `NFkB_KN` and $\mathcal{K}'$ = `NFkB_KNN` are the total kinase and IKK pools.

$$\frac{dK_a}{dt} = k_a C_1 \left(\mathcal{K} - K_a\right) \frac{k_{A20}}{k_{A20} + A} - k_i K_a$$

$$\frac{dK_n}{dt} = k_4 \left(\mathcal{K}' - K_n - K^+ - K^-\right) - k_1 K_a^2 K_n$$

$$\frac{dK^+}{dt} = k_1 K_a^2 K_n - k_3 K^+ \frac{k_2 + A}{k_2}$$

$$\frac{dK^-}{dt} = k_3 K^+ \frac{k_2 + A}{k_2} - k_4 K^-$$

The activation input C_1 is bound surface TNFR1. The apoptosis test reads the internalized pool instead (M3.7). A20 enters twice: as a saturable attenuation of kinase activation and as a linear acceleration of IKK inactivation.

IκB and NF-κB shuttling

P is phosphorylated IκB; N_c and N_n cytoplasmic and nuclear NF-κB; I and I_n cytoplasmic and nuclear IκB; X and X_n the cytoplasmic and nuclear complexes; X_P the complex with phosphorylated IκB. K_V = 5 is the cytoplasm-to-nucleus volume ratio, on every transfer into a nuclear complex.

$$\frac{dP}{dt} = a_2 K^+ I - t_p P \qquad \frac{dX_P}{dt} = a_3 K^+ X - t_p X_P$$

$$\frac{dN_c}{dt} = c_{6a} X - a_1 N_c I + t_p X_P - i_1 N_c \qquad \frac{dN_n}{dt} = i_1 N_c - a_1 K_V I_n N_n$$

$$\frac{dI}{dt} = -a_2 K^+ I - a_1 I N_c + c_4 I_t - c_{5a} I - i_{1a} I + e_{1a} I_n$$

$$\frac{dI_n}{dt} = -a_1 K_V I_n N_n + i_{1a} I - e_{1a} I_n$$

$$\frac{dX}{dt} = a_1 I N_c - c_{6a} X - a_3 K^+ X + e_{2a} X_n \qquad \frac{dX_n}{dt} = a_1 K_V I_n N_n - e_{2a} X_n$$

Gene states and feedback

Three gene-occupancy variables lie on [0, 2], for two alleles:

$$\frac{dG_{A20}}{dt} = q_1 N_n (2 - G_{A20}) - q_2 I_n G_{A20} \qquad \frac{dG_{I\kappa B}}{dt} = q_1 N_n (2 - G_{I\kappa B}) - q_2 I_n G_{I\kappa B}$$

$$\frac{dG_R}{dt} = q_{1r} N_n (2 - G_R) - (q_{2rr} + q_{2r} I_n) G_R$$

$$\frac{dA_t}{dt} = c_1 G_{A20} - c_3 A_t, \qquad \frac{dA}{dt} = c_4 A_t - c_5 A, \qquad \frac{dI_t}{dt} = c_1 G_{I\kappa B} - c_3 I_t$$

The A20 and IκB genes share their constants and track one another from equal initial values. The reporter gene G_R has its own constants and a spontaneous detachment term, and it alone drives the outputs. IκB translation, c_4 I_t, appears in the cytoplasmic IκB equation.

Output pairs

c*{1rr}^{ct}*, c{1r}^{c} and c_{1r}^{t} are the per-agent transcription coefficients set in secretion ("Transcription coefficients").

$$\frac{d\,\mathrm{chem}_t}{dt} = c_{1rr}^{\mathrm{ct}} + c_{1r}^{\mathrm{c}} G_R - c_{3r}^{\mathrm{chem}} \mathrm{chem}_t \qquad \frac{d\,\mathrm{chem}}{dt} = c_4^{\mathrm{chem}} \mathrm{chem}_t - c_5^{\mathrm{chem}} \mathrm{chem} - e_3^{\mathrm{chem}} \mathrm{chem}$$

$$\frac{d\,\mathrm{TNF}_t}{dt} = c_{1rr}^{\mathrm{ct}} + c_{1r}^{\mathrm{t}} G_R - c_{3r}^{\mathrm{TNF}} \mathrm{TNF}_t \qquad \frac{d\,\mathrm{TNF}_{\mathrm{int}}}{dt} = c_4^{\mathrm{TNF}} f_{\mathrm{dep}} \mathrm{TNF}_t - c_5^{\mathrm{TNF}} \mathrm{TNF}_{\mathrm{int}} - e_3^{\mathrm{TNF}} \mathrm{TNF}_{\mathrm{int}}$$

$$\frac{d\,\mathrm{ACT}_t}{dt} = c_{1rr}^{\mathrm{ACT}} + c_{1r} G_R - c_{3r}^{\mathrm{ACT}} \mathrm{ACT}_t \qquad \frac{d\,\mathrm{ACT}}{dt} = c_4^{\mathrm{ACT}} \mathrm{ACT}_t - c_5^{\mathrm{ACT}} \mathrm{ACT}$$

$$\frac{d\,\mathrm{IAP}_t}{dt} = c_{1rr}^{\mathrm{IAP}} + c_{1r} G_R - c_{3r}^{\mathrm{IAP}} \mathrm{IAP}_t \qquad \frac{d\,\mathrm{IAP}}{dt} = c_4^{\mathrm{IAP}} \mathrm{IAP}_t - c_5^{\mathrm{IAP}} \mathrm{IAP}$$

The chemokine and TNF pairs take per-agent, per-state coefficients and carry an export term. The chemokine export is applied on copy-back ("Chemokine and TNF output"); the TNF export drives membrane TNF. The activation and apoptosis-inhibitor pairs take their coefficients from parameters, carry no export, and are read by the decisions of M3.7.

f_dep is 1 normally, 0 under full TNF depletion, and `Core_depletionFractionTNF` under fractional depletion. It is the block's only depletion coupling (M8).

On copy-back the normalized activation and apoptosis-inhibition levels are formed as $\widehat{\mathrm{ACT}} = \mathrm{ACT} \cdot$ `NFkB_c3rACT` and $\widehat{\mathrm{IAP}} = \mathrm{IAP} \cdot$ `NFkB_c3rIAP`. The decisions read these, recomputed at every copy-back; they hold no history.

Transcription coefficients

The three per-agent coefficients are set in the secretion half-step by the NF-κB secretion helper, from three state-dependent arguments and an infection-bonus term. It is called once per macrophage per half-step under NF-κB dynamics, with values selected by state, and for no other cell type.

Macrophage state	Chemokine, reporter-dependent	TNF, reporter-dependent	Shared, reporter-independent
Resting	$\varepsilon_2 c_{1r}$	$\varepsilon_2 c_{1r}$	0
Infected	c_{1r}	c_{1r}	$\varepsilon_1 c_{1r}$
Chronically infected	c_{1r}	c_{1r}	c_{1r}
Activated	c_{1r}	c_{1r}	c_{1r}
Suppressed, any state	0	0	0
Dead	Not called	Not called	Not called

Here c_{1r} = `NFkB_c1r`, ε_2 = `NFkB_epsilon2` (resting to infected reporter-dependent ratio) and ε_1 = `NFkB_epsilon1` (reporter-independent to reporter-dependent ratio). The arguments give a resting macrophage a reduced reporter-driven rate and no basal part; infection adds a basal part and the full reporter rate; chronic infection or activation makes the basal part full. The suppressed path precedes the state switch and returns, so a suppressed macrophage transcribes neither species.

The helper assigns the two reporter-dependent coefficients as $x \cdot (bx)$ and the shared one as $x + bx$, where b is the infection bonus. Every call site passes b = 0. Both reporter-dependent coefficients are therefore zero for every state, and the shared coefficient equals its argument. By default the helper is reached only on the suppressed path, with every argument zero, so all three assignments give zero.

Macrophage state	Effective chemokine and TNF transcript source
Resting	0

Macrophage state	Effective chemokine and TNF transcript source
Infected	$\varepsilon_1 c_{1r}$
Chronically infected, activated	c_{1r}
Suppressed	0

Under NF-κB dynamics the reporter-dependent transcription coefficients are zero in every state, so macrophage chemokine and TNF output does not follow reporter-gene activity. Further consequences:

- gene states, feedback and the activation and apoptosis-inhibitor outputs are unaffected and still reach decisions (M3.7);
- resting macrophages produce neither species, whereas a coarse-grained resting NF-κB-positive macrophage secretes both at half rate; with `NFkB_epsilon2` = 0 the resting arguments are also zero, so the assignment changes resting behaviour only if $\varepsilon_2 > 0$;
- chronically infected and activated macrophages produce these species identically;
- at full reporter occupancy (G_R = 2) the source for a chronically infected macrophage is 0.05.

Chemokine and TNF output

On copy-back, after all subdivisions of the substep, three increments are applied at the agent's own compartment, from the post-integration intracellular chemokine and the full molecular substep:

$$\Delta C_{\mathrm{CCL2}} = e_3^{\mathrm{chem}}\,\mathrm{chem}\,\Delta t_m, \qquad \Delta C_{\mathrm{CCL5}} = e_3^{\mathrm{chem}}\,\mathrm{chem}\,\Delta t_m, \qquad \Delta C_{\mathrm{CXCL9}} = 2\,e_3^{\mathrm{chem}}\,\mathrm{chem}\,\Delta t_m.$$

Subdivision refines the integration but does not multiply the export. The writes are increments, not assignments; no block reads a chemokine field, so the chemokine fields carry no ownership constraint, and caseum keeps writing them.

The export is not conserved. The equation loses $e_3 \int \mathrm{chem}\,dt$ over the substep, while the grid gains $e_3\,\mathrm{chem}(\Delta t_m)\,\Delta t_m$; the two agree only if the intracellular pool is stationary. The TNF export is conserved, both its terms being inside the integrated system.

CCL5 and CXCL9 are reassigned from CCL2 after each solve, using ratios recomputed from the macrophage secretion rates on every call, so the block's direct writes to them are replaced, except those of the final half-step, which agents read (S6.14). The block writes after each molecular substep in a literal 1 : 1 : 2, `NFkB_e3Chem` $\times\mathrm{chem} \times \Delta t_m$ to CCL2 and CCL5 and twice that to CXCL9. This matches the default ratios but is not recomputed from them. The NF-κB secretion helper writes rate variables only. S6.14 holds the derived-chemokine relation; under this block it holds while the three macrophage chemokine rates keep the default ratios, which the literal equals.

Caseum secretion is a direct field increment at a quarter of the macrophage chemokine rates and does not pass through the macrophage secretion routine (S5.6). Under this block a macrophage writes no chemokine during secretion, its chemokine arriving through the copy-back increments, so the quarter relationship between caseum and macrophage output no longer holds.

What is not secreted

Under NF-κB dynamics the macrophage secretion routine does not call the chemokine helper or the TNF helper; the NF-κB helper is called in their place.

- The copy-back increments are the only macrophage chemokine source. The IL-10 step on chemokine secretion (S6.10) has no counterpart, so IL-10 does not inhibit macrophage chemokine production. With the absent IL-10 coupling to TNF production, IL-10 then acts on macrophages only through STAT3 (M3.7) and Treg contact (S3.10).
- The TNF helper's rate variables keep the values of any earlier path, zero for a macrophage created under this configuration; this has no effect, the transcript being frozen. Macrophage TNF reaches the field only through export to membrane TNF and its release by the TNFR block.
- IL-10 secretion is unchanged (M3.4).

T cells get no NF-κB block and call the TNF helper whatever the flag, its increment gated on the simulation's TNF receptor flag. Under NF-κB dynamics alone, T cells therefore secrete TNF by coarse-grained increment while macrophages use the receptor chain, on one field (M3.8).

Equilibration at construction

Under NF-κB dynamics every macrophage equilibrates its pathway in its constructor, before placement and before any timestep, so that it starts near a resting steady state without TNF. The loop runs $n = 7200\nu$ iterations of $\delta t = 6/\nu$ s, a total of 12 h whatever ν, where ν = `NFkB_NFkBTimeCoefficient`.

The initial total NF-κB pool, drawn at construction (S6.3), is placed in the cytoplasmic complex; equilibration redistributes it. The routine is a separate implementation in which 14 of the 26 derivatives are held at zero and the rest are as written above.

Group	Species	Integrated
Receptor-proximal kinases	K_a, K_n, K^+, K^-	No
Phosphorylated forms	P, X_P	No
NF-κB shuttling	N_c, N_n, X, X_n	Yes

Group	Species	Integrated
IκB	I, I_n, I_t	Yes
A20	A, A_t	Yes
Gene states	G_A20, $G_{I\kappa B}$, G_R	Yes
Output pairs	All eight output species	No

Without TNF, the kinase and phosphorylation derivatives are zero at the initial condition ($K_a = K^+ = K^- = X_P = P = 0$, K_n = `NFkB_KNN`) and stay zero, so holding these six species is exact. The retained terms that involve them vanish.

Freezing the output pairs is a simplification, not exact: the reporter state is integrated and changes. A new macrophage starts with its output species at zero, and they rise over its first hours. The normalized levels computed at the end of the routine are therefore zero.

The update is a hand-written explicit forward Euler step, all increments taken from pre-update values. It is independent of the stepper and of `--odesolver`. The step derives from a fixed 6 s reference; it equals Δt_m only under the spectral diffusion option (M1). No convergence test is applied. A reimplementation must reproduce the 12 h and the held set, not necessarily the Euler scheme. No other cell type equilibrates.

The initial NF-κB pool is drawn as exp(N(m, s)), with m = `NFkB_meanNFkB` and s = 0.64872 m as the parameters of the underlying normal; with the default mean of 10^5, the drawn pool overflows to infinity for about 94 % of agents and is zero for the rest, which affects only NF-κB dynamics.

Fields written

Field	Mode	Default
CCL2, CCL5, CXCL9	Increment	Written coarsely (S6)

M3.6 MHC class II dynamics

Summary. MHC class II antigen presentation is integrated per macrophage in a 16-equation block, whose IFN-γ and antigen inputs are imposed each substep from the counts of Tgam-holding compartments and extracellular bacteria nearby. A macrophage whose surface antigen-loaded MHC class II reaches a threshold moves every neighbouring cognate T cell in its T0 state to its active state. Up to four switchable factors let the macrophage's own bacterial burden suppress the pathway.

Notation and units

Ligands and peptides are molar: free IFN-γ g, endosomal antigen a, endosomal antigen-derived peptide q and endosomal self-peptide s. Receptors and MHC species are molecules per cell: free IFN-γ receptor $\mathcal{R}$, receptor-ligand complex $\mathcal{C}$, CIITA mRNA m, CIITA protein p, MHC class II mRNA μ, and six MHC species, H (free), S (self-loaded) and E (antigen-loaded), each endosomal (in) and surface (out).

A binding event written into a molar equation is divided by N_A and V_MIIC = `Molecular_MHCII_volMIIC`; count equations need no conversion. Receptor occupancy is $\omega = \mathcal{C}$ / `Molecular_MHCII_recPerCell`. The down-regulation factors d_1 to d_4 are 1 when their switches are clear.

The equations

IFN-γ receptor:

$$\frac{d\mathcal{R}}{dt} = -k^g_{\text{on}} g \mathcal{R} + k^g_{\text{off}} \mathcal{C} + k_{\text{recyc}} \mathcal{C} \qquad \frac{d\mathcal{C}}{dt} = k^g_{\text{on}} g \mathcal{R} - k^g_{\text{off}} \mathcal{C} - k_{\text{recyc}} \mathcal{C}$$

The receptor total is conserved: recycling returns the receptor in one step, and there is no synthesis or degradation. A cell's IFN-γ receptor complement is fixed for life, unlike its TNF and IL-10 receptor complements. g is imposed ("Inputs"), so ligand bound is debited from no pool.

Transcription and translation, with σ_m = `Molecular_MHCII_mscale`:

$$\frac{dm}{dt} = k^M_{\text{txn}} (1 + \sigma_m \omega) - k^M_{\text{deg}} m \qquad \frac{dp}{dt} = k^P_{\text{tsl}} m \, d_2 - k^P_{\text{deg}} p \qquad \frac{d\mu}{dt} = k_{\text{txn}} p \, d_4 - k^\mu_{\text{deg}} \mu$$

Occupancy amplifies CIITA transcription by $1 + \sigma_m \omega$ above a constitutive floor, without saturating.

Antigen and peptide:

$$\frac{da}{dt} = \frac{k_{\text{pino}} \, a_{\text{out}}}{V_{\text{MIIC}}} - k^a_{\text{deg}} a \, d_1 - k_{\text{toLys}} a$$

$$\frac{dq}{dt} = \frac{k_{\text{pino}} \, q_{\text{out}}}{V_{\text{MIIC}}} + k^a_{\text{deg}} a \, d_1 + \frac{k^H_{\text{off}} E_{\text{in}} d_3 - k^H_{\text{on}} q H_{\text{in}}}{N_A V_{\text{MIIC}}} - k_{\text{toLys}} q$$

Antigen is taken up by pinocytosis from the imposed a_out, then processed into peptide or routed to the lysosome. Processing is an exact gain to peptide, and d_1 scales both sides. A second peptide source is pinocytosis of preformed peptide at the fixed concentration q_out = `Molecular_MHCII_exPepOut`, independent of burden; at zero, antigen is the only source.

Self-peptide:

$$\frac{ds}{dt} = k_{\text{source}} + \frac{k^H_{\text{deg}}\,(S_{\text{in}} + S_{\text{out}}) - k^H_{\text{on}} H_{\text{in}} s + k^H_{\text{off}} S_{\text{in}} d_3}{N_A V_{\text{MIIC}}} - k_{\text{toLys}} s$$

Degraded self-loaded complexes, endosomal and surface, replenish endosomal self-peptide. Antigen-derived peptide has no such replenishment, so the two peptides are treated asymmetrically.

MHC class II species, with σ_H = `Molecular_MHCII_mhcScale`:

$$\frac{dH_{\text{in}}}{dt} = k_{\text{tsl}}\,(1 + \sigma_H \omega)\,\mu\, d_2 - k^H_{\text{on}} H_{\text{in}} s + k^H_{\text{off}} S_{\text{in}} d_3 - k^H_{\text{on}} H_{\text{in}} q + k^H_{\text{off}} E_{\text{in}} d_3 - k_{\text{out}} H_{\text{in}} + k_{\text{in}} H_{\text{out}} - k^H_{\text{deg}}\,(1 - \omega)\, H_{\text{in}}$$

$$\frac{dH_{\text{out}}}{dt} = k_{\text{out}} H_{\text{in}} - k_{\text{in}} H_{\text{out}} - k^H_{\text{deg}} H_{\text{out}}$$

$$\frac{dS_{\text{in}}}{dt} = k^H_{\text{on}} H_{\text{in}} s - k^H_{\text{off}} S_{\text{in}} d_3 - k_{\text{out}} S_{\text{in}} + k_{\text{in}} S_{\text{out}} - k^H_{\text{deg}} S_{\text{in}} \qquad \frac{dS_{\text{out}}}{dt} = k_{\text{out}} S_{\text{in}} - k_{\text{in}} S_{\text{out}} - k^H_{\text{deg}} S_{\text{out}}$$

$$\frac{dE_{\text{in}}}{dt} = k^H_{\text{on}} H_{\text{in}} q - k^H_{\text{off}} E_{\text{in}} d_3 - k_{\text{out}} E_{\text{in}} + k_{\text{in}} E_{\text{out}} - k^H_{\text{deg}} E_{\text{in}} \qquad \frac{dE_{\text{out}}}{dt} = k_{\text{out}} E_{\text{in}} - k_{\text{in}} E_{\text{out}} - k^H_{\text{deg}} E_{\text{out}}$$

Each loading state exists in an endosomal and a surface form, exchanged at k_out and k_in and degraded at k_deg^H; loading occurs only in the endosome. E_out, surface antigen-loaded MHC class II, is the block's output.

New MHC class II enters H_in only, amplified by occupancy and scaled by d_2, which also scales CIITA translation. H_in degradation carries the factor $(1 - \omega)$, so receptor engagement spares free endosomal MHC class II; every other species degrades at the plain rate. The factor is unbounded below, but $\omega \leq 1$ holds provided `recPerCell` is consistent with the initial condition. Every unloading term carries d_3; the self-peptide replenishment term carries the plain degradation rate.

Inputs

Free IFN-γ and external antigen are not fields; neither is represented as a field anywhere in the model (spine 4.1). Before each integration they are computed from the macrophage's neighbourhood and written into their vector positions, overwriting the previous values; on copy-back the members hold the imposed values.

The inputs are: compartments within Moore distance 1, the centre included, holding at least one Tgam of any state, including dead Tgam not yet removed (n_1); compartments at distance exactly 2 holding at least one Tgam (n_2); and live plus dead extracellular bacteria summed over distance 1 (n_Mtb). All use periodic wrapping. The same function supplies the MHC class II input statistics (S9.15).

$$g = f_1 n_1 + f_2 n_2, \qquad a_{\text{out}} = f_a n_{\text{Mtb}},$$

with f_1 = `Molecular_MHCII_tgamInMN1ToMolarConcIFNg`, f_2 = `Molecular_MHCII_tgamInMN2ToMolarConcIFNg` and f_a = `Molecular_MHCII_mtbToMolarConcAntigen`, which convert counts to molar concentrations. The two IFN-γ factors weight distance. The inputs count compartments, so two Tgam in one compartment count once.

Under this component a macrophage processes antigen in proportion to the local bacterial count and, once its surface complex count crosses a threshold, moves neighbouring T0 T cells to their active state. Proliferation still requires an exposed macrophage (S3).

This is a second Tgam-to-macrophage pathway, added to the contact-mediated IFN-γ action of S3.7. It counts Tgam regardless of producer status and acts on MHC class II, not STAT1, so the two pathways count different sets of cells.

The derivative routine writes 14 of the 16 entries; no build assigns the derivatives of the two imposed positions. Within a substep the integrator advances them with whatever the shared derivative buffer holds (M3.2), and the binding and uptake terms read them at every stage of a multi-stage step. They are refreshed at the next substep, so nothing accumulates across substeps; under NF-κB subdivision they advance unrefreshed across all subdivisions of the substep.

An alternative formulation, in which the two inputs are integrated state in a finite reaction volume, is present but cannot be reached from any parameter file or command line. `Molecular_MHCII_numCells` and `rxnVol` are read by no live code; `kdegIfng` is rescaled at load time but consumed only by that formulation; `ifng` and `exAgOut` are read at construction as initial values and overwritten by the first imposition.

T cell activation

The block's output predicate is

$$\text{canActivate} \iff E_{\text{out}} \geq \texttt{Molecular_MHCII_mhcIIBindingThreshold},$$

a plain non-strict comparison, without the seven-figure truncation and without a draw.

When the component is selected, a macrophage performs the activation action in its next-state computation before the suppression test, unless it is chronically infected. If the predicate holds, every cognate T cell within radius r = `Molecular_MHCII_tcellActivationRadius` is activated. With r = 0 only the macrophage's own compartment is visited; with $r \geq 1$, $(2r + 1)^2$ compartments of the periodic Moore neighbourhood. Every slot is examined, and a cognate T cell of any type is activated; non-cognate cells are skipped. There is no probability, no per-timestep limit and no record, so repeat activations in one pass have no extra effect.

On the receiving side, a T cell in its T0 state has both its current and next state set to active, and its per-type activation counter is incremented; in any other state nothing happens. Writing the current state is an immediate cross-agent effect, visible to agents processed later in the pass (spine 2.3). It bypasses the guarded next-state setter, so a T0 cell killed earlier in the timestep is returned to the active state. The Treg suppression routine guards against this; no T0 cell exists by default.

Recruited T cells are constructed in their T0 state when this component is enabled (S8.12), and active otherwise, under either recruitment option. Initial cognate Tgam (zero by default) are created active whatever the setting. The T0 states are therefore unreachable in the default configuration, and this component is the only route out of them.

The program help text says this component requires the lymph-node recruitment option; the program does not require it, and under probabilistic recruitment recruited T cells start in T0.

Observation

The MHC class II input averages, the mean surface antigen-loaded MHC class II and the count of macrophages satisfying the predicate are computed every timestep in every configuration (S9.15); they are zero or inert by default. Further outputs are written only when this component is selected: the statistics comparing threshold-satisfying macrophages with T cells by distance from the centre of mass of the current timestep's granuloma region, and the per-type T cell activation counts (M3.6, "T cell activation"). Both are zero by default.

Down-regulation by intracellular bacteria

Four factors let a macrophage's live intracellular bacterial count C suppress its own presentation:

$$d_i = 1 - \frac{r_i C}{K_i + C},$$

with r_i the maximum fractional suppression and K_i the half-saturation burden: 1 at C = 0, $1 - r_i$ at saturation and $1 - r_i/2$ at C = K_i. Each has its own switch and d_i = 1 when it is clear; any combination may be enabled. Each is computed once per agent per substep and held across the stages of a step.

Factor	Parameters	Terms scaled	Occurrences
d_1	`mhcMtbDownReg1`, `mtbDownRegRate1`, `mtbHalfDownReg1`	Antigen processing, k_deg^a	2
d_2	`mhcMtbDownReg2`, `mtbDownRegRate2`, `mtbHalfDownReg2`	CIITA translation and MHC class II translation	2
d_3	`mhcMtbDownReg3`, `mtbDownRegRate3`, `mtbHalfDownReg3`	Peptide unloading, k_off^H	5
d_4	`mhcMtbDownReg4`, `mtbDownRegRate4`, `mtbHalfDownReg4`	MHC class II transcription, k_txn	1

All are in group `Molecular_MHCII`. Each factor scales every occurrence of its rate: d_1 diverts antigen to the lysosome rather than destroying it, and d_3 keeps unloading mass-conserving. d_2 acts twice in the transduction chain and d_4 once, so at steady state d_2 enters quadratically.

d_3 scales the dissociation rate; the association rate is never scaled. d_3 therefore raises surface antigen-loaded MHC class II with burden, in the opposite direction to the other three.

C is read through a per-class accessor that returns the live intracellular count for macrophages and zero for every class that does not override it, all T cells included, so the factors are 1 for T cells. A macrophage with no intracellular bacteria has all factors at 1. Chronically infected macrophages are excluded from the activation call, so for them the factors change only internal state. The factors can change a decision through infected macrophages and through activated macrophages that still hold bacteria; under the default parameter ranges either holds fewer than the chronic threshold of 10–15.

The rates r_i are not bounded above, so $d_i < 0$ once $C > K_i/(r_i - 1)$. With a switch set and K_i = 0, a macrophage with no intracellular bacteria evaluates 0/0, and the resulting not-a-number propagates until the copy-back check of M3.2 throws.

In the default configuration all four switches are clear and every rate and half-saturation is zero, so down-regulation is inactive even with the component selected, and neither a negative factor nor a not-a-number arises.

Rescaling at load time

Thirteen `Molecular_MHCII` rate constants are declared per 5 s, and `konIfng` per hour. At load time, after the diffusion method is set, each is multiplied once by Δt_m / 5 s, which is 1.2 under the selected spectral option, 1.0 under the alternating-direction explicit option and 0.6 under Crank–Nicolson ADI (M1). The rescaling runs unconditionally, whether or not the component is selected. The rescaled values are written in place in the parameter store, so a repeated call rescales them again.

On the restart path the rescaling can run more than once, each run multiplying the rates by a further Δt_m / 5 s ($1.2^2 = 1.44$ times their file values after two runs); the rescaled values are unused by default.

`Molecular_MHCII_konIfng`, declared per molar per hour, is rescaled by the same factor as the thirteen rates declared per 5 s.

The rescaling converts a per-5 s rate to a per-Δt_m rate.

Quantities derived at load time

Thirteen quantities are derived at load time from the rescaled rates and the composition parameters, and any parameter-file value for them is overwritten. The file gives zero for the six derived rates and omits the seven derived initial conditions. The program does not guard the derivations: zero `m`, `p` or `mrnaPerCell`, or `propMhcIn` = 1, divides by zero, with no consequence by default.

Six rates, in dependency order, with π_{in} = `propMhcIn`, μ_0 = `mrnaPerCell` and s_0 the derived initial self-peptide:

$$k^M_{\text{txn}} = k^M_{\text{deg}} m \qquad k^P_{\text{tsl}} = k^P_{\text{deg}} p/m \qquad k_{\text{in}} = \frac{\pi_{\text{in}}}{1 - \pi_{\text{in}}} k_{\text{out}} - k^H_{\text{deg}}$$

$$k_{\mathrm{txn}} = k^{\mu}_{\mathrm{deg}}\mu_0/p \qquad k_{\mathrm{tsl}} = k^{H}_{\mathrm{deg}}\frac{H_{\mathrm{in}} + H_{\mathrm{out}} + S_{\mathrm{in}} + S_{\mathrm{out}}}{\mu_0} \qquad k_{\mathrm{source}} = k_{\mathrm{toLys}}\, s_0$$

The in-force values are $k^{M}_{\mathrm{txn}} = 1.2695 \times 10^{-4}$, $k^{P}_{\mathrm{tsl}} = 3.3250 \times 10^{-3}$, $k_{\mathrm{in}} = 1.1315 \times 10^{-3}$, $k_{\mathrm{txn}} = 9.9260$, $k_{\mathrm{tsl}} = 1.1113 \times 10^{-4}$ and $k_{\mathrm{source}} = 3.8538 \times 10^{-6}$.

Seven initial conditions:

Quantity	Value	Species
`recOut`	10^4	Free IFN-γ receptors, set to `recPerCell`
`mrna`	10^5	MHC class II mRNA
`mhcIn`	1.3333×10^4	Free endosomal MHC class II
`mhcOut`	2.6667×10^4	Free surface MHC class II
`mhcSelfPepIn`	5.3333×10^4	Endosomal self-peptide complexes
`mhcSelfPepOut`	1.0667×10^5	Surface self-peptide complexes
`selfPepIn`	4.8166×10^{-4} M	Endosomal self-peptide

The four MHC initial values sum to `mhcPerCell` = 2×10^5, with loaded share `propMhcBound` = 0.8 and endosomal share `propMhcIn` = 1/3. The two antigen-loaded species start at zero and are not derived, so a macrophage starts with no antigen-loaded MHC class II; the activation threshold of 150 is about three orders of magnitude below the surface self-loaded pool. k_out / k_in = 2.098; the degradation term in the derivation of k_in makes it differ from $(1 - \pi_{\mathrm{in}})/\pi_{\mathrm{in}}$ = 2.

Occupancy at the imposed exposures

$f_1/f_2 = 3.555$, so a Tgam-holding compartment at distance 1 contributes about 3.5 times as much IFN-γ as one at distance 2. At the in-force rates one Tgam-holding compartment at distance 1 gives a quasi-steady $\mathcal{C}/\mathcal{R} = 244$ and $\omega = 0.996$. The CIITA transcription amplification is then $1 + \sigma_m\omega = 45.7$, the MHC class II translation amplification $1 + \sigma_H\omega = 5.16$, and H_in degradation runs at 0.4 % of its nominal rate. The response to local Tgam density is compressed below one adjacent Tgam-holding compartment and is effectively binary.

M3.7 Switched decision points

Five decisions read a quantity that depends on the selection. In each, the decision's threshold-and-probability form and its place in the schedule are unchanged; what changes is the quantity substituted for the coarse-grained occupancy proxy and, in three cases, the parameter group.

Decision	Coarse-grained quantity (S6)	Fine-grained quantity	Parameter group switches
TNF-induced apoptosis	Estimated internalized bound TNFR1 (S6.6)	Integrated I_1; IAP-modulated rate under NF-κB	Yes
Macrophage NF-κB activation	Signalling quantity (S6.5)	Integrated C_1; $\widehat{\mathrm{ACT}}$ under NF-κB	Yes
TNF signalling quantity (helper)	Estimated bound TNFR1, saturated against free receptor (S6.4)	Integrated C_1, saturated against total receptor	No
STAT3 activation	Estimated bound IL-10R (S6.7)	Integrated B	Yes
Treg suppression strength	Producer test on `Core_Tcell_Treg_dIL10`	Producer test on `Molecular_IL10_Tcell_IkSynthTcell`	No

Treg suppression strength is evaluated on the Treg class; the other decisions are on the base agent class. All take the flags as call arguments, so each uses the flags passed at its call site.

TNF-induced apoptosis

Three branches, chosen on the NF-κB flag first, then the TNF receptor flag. The coarse-grained branch is specified in S6.6.

TNF receptor dynamics without NF-κB. The quantity is the integrated internalized TNFR1 count I_1, with the `Molecular_TNF` threshold, rate and ceiling. The ceiling is applied to a separate variable: the threshold test reads unclamped I_1, the probability the clamped value.

NF-κB dynamics. The quantity is I_1, unclamped in both tests, against `Molecular_TNF_thresholdApoptosisTNF`, and the rate constant is

$$k^{\mathrm{eff}} = \mathtt{NFkB_kApoptosis} \cdot \frac{\mathtt{NFkB_kIAP}}{\mathtt{NFkB_kIAP} + \widehat{\mathrm{IAP}}}.$$

This is one of the two routes from pathway outputs to decisions, the other being macrophage NF-κB activation. Both survive the zero transcription coefficients of M3.5.

All three T cell types call this test with the NF-κB argument fixed false, not the simulation flag. They take the TNF receptor branch when that component is on and the coarse-grained branch otherwise. Under NF-κB dynamics alone, macrophages test apoptosis on the integrated pool and T cells on the coarse-grained estimate, from one field.

Macrophage NF-κB activation

Three branches on the same flags, reached only from the macrophage pathway-stimulation routine. The coarse-grained branch is specified in S6.5.

TNF receptor dynamics without NF-κB. The threshold test reads the raw integrated C_1 against `Molecular_TNF_Mac_thresholdNFkBTNF`; the probability reads the TNF signalling quantity, with `Molecular_TNF_Mac_kNFkB`.

NF-κB dynamics. Both tests read $\widehat{\mathrm{ACT}}$ against `NFkB_actThreshold` and `NFkB_activationRate`. The signalling quantity is still computed on entry and discarded.

TNF signalling quantity

The helper returns a quantity of TNF signalling in bound receptors per cell. It takes the TNF receptor flag only. With the flag on:

$$\min(C_1,\ \phi_{\mathrm{sat}}\,(R_1 + C_1))\,, \qquad \phi_{\mathrm{sat}} = \texttt{Molecular_TNF_Mac_saturationFracNFkBTNF}.$$

With the flag off it is the coarse-grained formula of S6.4, saturated with `Core_Mac_saturationFracNFkBTNF` against free R_1.

The fraction applies to the total pool in the fine-grained branch and to the free pool in the coarse-grained one. The two coincide while the pool is fixed; under TNF receptor dynamics the free pool is depleted by binding, and the total keeps the ceiling stable. A reimplementation must not unify them. In the fine-grained branch the quantity is an integrated state with full history, not an equilibrium value recomputed from the current field.

The helper's parameters come from macrophage groups although it is on the base class. Its callers are the macrophage activation test and the M4 TGF-β activation sweep. Its coarse-grained branch reads the agent's own free TNFR1, so receptor heterogeneity enters both branches.

STAT3 activation

The coarse-grained branch, with no ceiling, is specified in S6.7. Under IL-10 receptor dynamics the quantity is the integrated bound IL-10R count B, with the `Molecular_IL10_Mac` threshold and rate, and a ceiling appears:

$$\min(B,\ \psi_{\mathrm{sat}}\,(B + R_{\text{IL-10}}))\,, \qquad \psi_{\mathrm{sat}} = \texttt{Molecular_IL10_Mac_saturationFracSTAT3IL10}.$$

The ceiling is on a separate variable read by the probability only; the threshold reads raw B.

Treg suppression strength

The Treg's producer test, which selects the suppression probability (S3.10), reads `Core_Tcell_Treg_dIL10` $\neq 0$ when the IL-10 receptor flag is clear and `Molecular_IL10_Tcell_IkSynthTcell` $\neq 0$ when it is set. It is a parameter test, fixed for the run.

Ceilings below thresholds

In three decisions, and only in their fine-grained branches, the ceiling bounds the probability argument but not the threshold test. If a ceiling lies below its threshold, the threshold can pass while the probability argument is negative, so the event never fires at any signal level. The conditions are

$$\texttt{Molecular_TNF_saturationApoptosisTNF} < \texttt{Molecular_TNF_thresholdApoptosisTNF}, \qquad \psi_{\mathrm{sat}}\,(B + R_{\text{IL-10}}) < \texttt{Molecular_IL10_Mac_thresholdSTAT3IL10},$$

and analogously for activation. The program does not check the ordering, and the STAT3 condition depends on per-agent receptor counts, so it can hold for some agents and not others. The program neither clamps the probability argument at zero nor declares the constraint. The coarse-grained branches are not exposed, since the apoptosis branch clamps before both tests and the STAT3 branch has no ceiling. At the default values the ordering holds at mean complements, by factors of 2.3 (apoptosis), about 10 (activation) and about 75 (STAT3).

M3.8 Partial enablement

NF-κB dynamics without TNF receptor dynamics integrates the TNFR equations through the failsafe (M3.2), in their NF-κB variant (M3.5), while the simulation's TNF receptor flag stays false and every other part of the program reads that flag.

Under this selection bound TNFR1 is integrated and drives the pathway, and membrane TNF accumulates and is released. The apoptosis and activation tests take their NF-κB branches, chosen on the NF-κB flag. The activated-macrophage IL-10 plasticity factor (M3.4) is non-zero.

Because the TNF receptor flag stays false:

- **TNF is consumed by both the block and the sink.** The coarse-grained TNF sink, gated on the clear flag, is applied while the block internalizes TNF. The integration loop runs before the sink loop, so the sink's decrement is applied after the block's assignment and survives into the next substep.
- **The sink reads a depleting pool.** Its occupancy proxy reads the free TNFR1 and TNFR2 counts that the block is integrating and depleting.
- **Depletion leaves consumption intact.** Under TNF depletion production stops but both consumption routes remain, the sink being exempt from depletion (M3.1).
- **The helper estimates coarsely from integrated counts.** Called with the clear flag, the signalling quantity takes its coarse-grained branch on the depleting integrated receptors. The activation test discards it, but the M4 TGF-β activation sweep uses it.
- **The two cell populations use different resolutions on one field.** T cells receive no block (M3.2).

	Macrophages	**T cells**
TNF production	Pathway export to membrane TNF, then release	Coarse-grained increment, gated on the clear flag
TNF consumption	Block internalization and coarse-grained sink	Coarse-grained sink only

	Macrophages	T cells
Apoptosis	Integrated internalized TNFR1, IAP-modulated	Coarse-grained estimate
Field write	Assignment	Increment

In a shared compartment the macrophage's assignment overwrites the T cell's earlier secretion increment, secretion preceding the substep loop.

Components	Status
TNFR alone	Flag and integrated blocks agree
IL-10R alone	Flag and integrated blocks agree; activated macrophages synthesize no IL-10 (M3.4)
TNFR and IL-10R	Flag and integrated blocks agree
MHC class II alone, or with either receptor component	Flag and integrated blocks agree; T cells receive no MHC class II block
NF-κB without TNFR	TNFR block integrated with the flag false, as described
NF-κB with TNFR	Flag and integrated blocks agree; the failsafe changes nothing

These effects arise exactly in the selections with NF-κB dynamics and without TNF receptor dynamics, whatever the other two components. Adding TNF receptor dynamics removes them.

M3.9 Parameters

Values are those in force in the default configuration. The module being off, most are read and never consumed. Status is given in the Meaning column: **active** (affects a trajectory or output by default), **read** (read by default but the quantity is never used), **module** (read only when its component is selected) or **inert** (read by no live code).

Only `Molecular_TNF_Mac_saturationFracNFkBTNF` and its neutrophil counterpart ([0, 1]) and the four down-regulation switches ([0, 1]) are range-checked. Every other parameter here, including every rate constant, threshold, pool and conversion factor, accepts any value, negative ones included (spine 7.0, "Validity ranges").

TNF receptor group

`Molecular_TNF_KD1` and `Molecular_TNF_KD2`, active at coarse resolution, are tabled in S6.16.

Identifier	Value	Units	Meaning
`Molecular_TNF_kOn1`	2.8×10^{7}	$M^{-1}\ s^{-1}$	TNFR1 association. Module
`Molecular_TNF_kOn2`	3.5×10^{7}	$M^{-1}\ s^{-1}$	TNFR2 association. Module
`Molecular_TNF_kInt1`	7.7×10^{-4}	s^{-1}	Internalization of bound TNFR1. Module
`Molecular_TNF_kInt2`	4.6×10^{-4}	s^{-1}	Internalization of bound TNFR2. Module
`Molecular_TNF_kRec1`, `kRec2`	1.8×10^{-5}	s^{-1}	Recycling of internalized receptor. Module
`Molecular_TNF_kDeg1`, `kDeg2`	5×10^{-5}	s^{-1}	Degradation of internalized receptor. Module
`Molecular_TNF_kShed`	5×10^{-4}	s^{-1}	Shedding of bound TNFR2. Module
`Molecular_TNF_kT1`, `kT2`	3.8×10^{-4}	s^{-1}	Receptor turnover; forms the per-agent synthesis rate at construction. Read
`Molecular_TNF_kTrans`	8×10^{-5}	s^{-1}	Translation of transcript to membrane TNF. Module
`Molecular_TNF_diffusivityShedTNFR2`	3.2×10^{-8}	$cm^2\ s^{-1}$	Shed TNFR2 diffusivity; consumed by M1. Module
`Molecular_TNF_Mac_kRNAMac`	1.5	s^{-1}	Ceiling of the transcription logistic, macrophages. Read each secretion half-step
`Molecular_TNF_Mac_kSynthMac`	0.21	molecules $cell^{-1}\ s^{-1}$	Floor of the transcription logistic, macrophages. Read each secretion half-step
`Molecular_TNF_Mac_kTaceMac`	4.4×10^{-4}	s^{-1}	Release of membrane TNF, macrophages. Read at construction
`Molecular_TNF_Tcell_kRNATcell`	0.15	s^{-1}	Ceiling, T cells. Read
`Molecular_TNF_Tcell_kSynthTcell`	0.021	molecules $cell^{-1}\ s^{-1}$	Floor, T cells; also the construction-time synthesis rate. Read
`Molecular_TNF_Tcell_kTaceTcell`	4.4×10^{-5}	s^{-1}	Release of membrane TNF, T cells. Read at construction
`Molecular_TNF_thresholdApoptosisTNF`	2028	molecules	Apoptosis threshold on I_1. Module
`Molecular_TNF_kApoptosis`	1.33×10^{-6}	s^{-1}	Apoptosis rate constant. Module
`Molecular_TNF_saturationApoptosisTNF`	4691	molecules	Apoptosis ceiling, probability argument only. Module
`Molecular_TNF_Mac_thresholdNFkBTNF`	81.35	molecules	Activation threshold on C_1. Module
`Molecular_TNF_Mac_kNFkB`	9.74×10^{-6}	s^{-1}	Activation rate constant. Module

Identifier	Value	Units	Meaning
`Molecular_TNF_Mac_saturationFracNFkBTNF`	0.541	fraction	Saturation fraction on the total receptor pool. Module

Derived: $k_{\mathrm{off},1} = 5.32 \times 10^{-4}$ s^{-1} and $k_{\mathrm{off},2} = 1.47 \times 10^{-2}$ s^{-1}, so bound TNFR2 dissociates about 28 times faster than bound TNFR1. The recycled fraction of internalized receptor is k_rec / (k_rec + k_deg) = 0.265 for both subtypes, so about three quarters is lost. The receptor pool relaxes with time constant 1/k_T = 2632 s, about 4.4 cellular timesteps.

IL-10 receptor group

`Molecular_IL10_IkD`, active at coarse resolution, is tabled in S6.16.

Identifier	Value	Units	Meaning
`Molecular_IL10_IkOn`	5.7×10^{5}	M^{-1} s^{-1}	Association. Module
`Molecular_IL10_IkT`	1.6×10^{-4}	s^{-1}	Turnover of free receptor; forms the synthesis rate at construction. Read
`Molecular_IL10_IkInt`	5×10^{-4}	s^{-1}	Internalization of bound receptor, a pure loss. Module
`Molecular_IL10_LinkRNAGamma`	20	bound receptors	Centre of the transcription logistic (M3.3). Module
`Molecular_IL10_LinkRNADelta`	7	bound receptors	Width of the same. Module
`Molecular_IL10_Mac_IkSynthMacInf`	0.02	molecules cell^{-1} s^{-1}	Synthesis argument, infected and chronically infected. Read each secretion half-step
`Molecular_IL10_Mac_IkSynthMacAct`	0.3	molecules cell^{-1} s^{-1}	Synthesis argument, activated. Read
`Molecular_IL10_Mac_halfSatIL10mRNA`	190	molecules	Half-saturation of the plasticity factor (M3.4). Read
`Molecular_IL10_Tcell_IkSynthTcell`	0.74	molecules cell^{-1} s^{-1}	Synthesis argument, all T cells; in effect the Treg rate. Read
`Molecular_IL10_Mac_thresholdSTAT3IL10`	10	molecules	STAT3 threshold on B. Module
`Molecular_IL10_Mac_kSTAT3IL10`	8×10^{-4}	s^{-1}	STAT3 rate constant. Module
`Molecular_IL10_Mac_saturationFracSTAT3IL10`	0.5	fraction	STAT3 ceiling, probability argument only. Module

Derived: $k_{\mathrm{off}} = 2.60 \times 10^{-4}$ s^{-1}. Of bound IL-10R, the fraction internalized rather than released is 0.66, each such event consuming ligand and receptor; the TNFR1 figure is 0.59, but about a quarter of that is recycled. The free pool relaxes with time constant 1/k_{l,T} = 6250 s, about 10.4 cellular timesteps.

Among the Treg suppression routines, only the unreachable five-argument routine and its callee read `IkD`, `IkOn`, `Molecular_TNF_KD1` and `kOn1` (S3.10); the live four-argument routine reads no receptor constant.

The receptor means and half-widths of both groups are active, owned by S6.3; their conversion to molecules per compartment is owned by S6.2.

NF-κB group

All entries are of module scope except `meanNFkB` and `KNN`.

Identifier	Value	Units	Meaning
`NFkB_KN`	10^{5}	molecules	Total IKK kinase pool
`NFkB_KNN`	2×10^{5}	molecules	Total IKK pool; also the declared initial neutral IKK of every agent. Read at construction
`NFkB_meanNFkB`	10^{5}	molecules	Location argument of the initial pool draw (S6.3, M3.5). Read at construction; gates a draw
`NFkB_ka`	2×10^{-6}	s^{-1}	Kinase activation per bound TNFR1
`NFkB_ki`	0.01	s^{-1}	Kinase inactivation
`NFkB_kA20`	10^{5}	molecules	Half-saturation of A20 attenuation
`NFkB_k1`	6×10^{-10}	s^{-1}	IKK activation, second order in active kinase
`NFkB_k2`	10^{4}	molecules	Michaelis constant of A20 acceleration of IKK inactivation
`NFkB_k3`	0.002	s^{-1}	IKK inactivation
`NFkB_k4`	0.001	s^{-1}	Recovery of inactive IKK
`NFkB_a1`	5×10^{-7}	s^{-1}	IκB association with NF-κB
`NFkB_a2`	10^{-7}	s^{-1}	IκB phosphorylation
`NFkB_a3`	5×10^{-7}	s^{-1}	Phosphorylation of complexed IκB
`NFkB_tp`	0.01	s^{-1}	Degradation of phosphorylated IκB
`NFkB_c5a`	10^{-4}	s^{-1}	Degradation of free IκB
`NFkB_c6a`	2×10^{-5}	s^{-1}	Degradation of complexed IκB
`NFkB_i1`	0.01	s^{-1}	Nuclear import of NF-κB
`NFkB_i1a`	0.002	s^{-1}	Nuclear import of IκB

Identifier	Value	Units	Meaning
`NFkB_e1a`	0.005	s^{-1}	Nuclear export of IκB
`NFkB_e2a`	0.05	s^{-1}	Nuclear export of the complex
`NFkB_q1`	4×10^{-7}	s^{-1}	NF-κB binding at the A20 and IκB promoters
`NFkB_q2`	10^{-6}	s^{-1}	IκB-induced detachment from those promoters
`NFkB_q1r`	10^{-7}	s^{-1}	NF-κB binding at the reporter
`NFkB_q2r`	10^{-7}	s^{-1}	IκB-induced detachment from the reporter
`NFkB_q2rr`	0.001	s^{-1}	Spontaneous detachment from the reporter
`NFkB_c1`	0.1	s^{-1}	A20 and IκB transcription
`NFkB_c3`	7.5×10^{-4}	s^{-1}	A20 and IκB transcript degradation
`NFkB_c4`	0.5	s^{-1}	A20 and IκB translation
`NFkB_c5`	5×10^{-4}	s^{-1}	A20 degradation
`NFkB_c1r`	0.05	s^{-1}	Reporter-driven transcription; supplied per state (M3.5)
`NFkB_epsilon1`	0.5	—	Reporter-independent to reporter-dependent ratio
`NFkB_epsilon2`	0	—	Resting to infected reporter-dependent ratio; scales a coefficient that is zero
`NFkB_c3rChem`, `c4Chem`, `c5Chem`, `e3Chem`	1.92×10^{-4}, 0.45, 5×10^{-4}, 1.39×10^{-5}	s^{-1}	Chemokine transcript degradation, translation, degradation, export
`NFkB_c3rTNF`, `c4TNF`, `c5TNF`, `e3TNF`	3.8×10^{-4}, 0.15, 5×10^{-4}, 2.49×10^{-6}	s^{-1}	TNF transcript degradation, translation, degradation, export to membrane
`NFkB_c1rrACT`, `c3rACT`, `c4ACT`, `c5ACT`	0.001, 1.92×10^{-4}, 0.5, 5×10^{-4}	s^{-1}	Activation-molecule pair; `c3rACT` also normalizes $\widehat{\mathrm{ACT}}$
`NFkB_c1rrIAP`, `c3rIAP`, `c4IAP`, `c5IAP`	0.001, 1.92×10^{-4}, 0.5, 5×10^{-4}	s^{-1}	Apoptosis-inhibitor pair; `c3rIAP` also normalizes $\widehat{\mathrm{IAP}}$
`NFkB_actThreshold`	24.8079	—	Activation threshold on $\widehat{\mathrm{ACT}}$
`NFkB_activationRate`	4.62441×10^{-3}	s^{-1}	Activation rate constant
`NFkB_kApoptosis`	1.38×10^{-6}	s^{-1}	Base apoptosis rate under NF-κB, slightly above `Molecular_TNF_kApoptosis` and just below the coarse-grained sampled range of S6.6; the threshold is `Molecular_TNF_thresholdApoptosisTNF`
`NFkB_kIAP`	38.5744	—	Half-inhibition constant of the apoptosis rate
`NFkB_NFkBTimeCoefficient`	6	—	Substep subdivisions (M3.2) and equilibration scaling (M3.5)

The equilibration also reads the pathway constants in its retained expressions, so enabling the component causes most of this group to be read at the construction of every macrophage. At $\nu = 6$ the equilibration runs 43 200 steps of 1 s, and each 6 s substep is integrated in six steps of 1 s.

The activation threshold corresponds to 1.29×10^5 raw activation molecules, and the IAP half-inhibition constant to 2.01×10^5. At the constitutive steady state without reporter activity, $\widehat{\mathrm{ACT}} = \widehat{\mathrm{IAP}} = 1.0$ (5.21×10^3 molecules): activation is pathway-driven, and at rest IAP lowers the apoptosis rate by 2.5 %.

MHC class II group

Rescaled at load time (M3.6), declared per 5 s except `konIfng`:

Identifier	Declared	In force	Units (declared)	Meaning
`Molecular_MHCII_konIfng`	5.9944×10^7	7.1933×10^7	$M^{-1}\,h^{-1}$	IFN-γ receptor association
`Molecular_MHCII_koffIfng`	4.8897×10^{-2}	5.8676×10^{-2}	per 5 s	IFN-γ receptor dissociation
`Molecular_MHCII_krecyc`	1.2265×10^{-3}	1.4718×10^{-3}	per 5 s	Receptor internalization and recycling
`Molecular_MHCII_kdegM`	1.0579×10^{-4}	1.2695×10^{-4}	per 5 s	CIITA mRNA degradation
`Molecular_MHCII_kdegP`	2.7708×10^{-3}	3.3250×10^{-3}	per 5 s	CIITA protein degradation
`Molecular_MHCII_kdegMrna`	8.2717×10^{-5}	9.9260×10^{-5}	per 5 s	MHC class II mRNA degradation
`Molecular_MHCII_kpino`	3.0416×10^{-16}	3.6499×10^{-16}	L per cell per 5 s	Pinocytic uptake
`Molecular_MHCII_kdegAg`	8.1837×10^{-3}	9.8204×10^{-3}	per 5 s	Antigen processing
`Molecular_MHCII_ktoLys`	6.6676×10^{-3}	8.0011×10^{-3}	per 5 s	Lysosomal routing
`Molecular_MHCII_konMhc`	5.5556×10^2	6.6667×10^2	M^{-1} per 5 s	Peptide loading
`Molecular_MHCII_koffMhc`	6.6759×10^{-2}	8.0111×10^{-2}	per 5 s	Peptide unloading
`Molecular_MHCII_kdegMhc`	4.6305×10^{-5}	5.5566×10^{-5}	per 5 s	MHC class II degradation

Identifier	Declared	In force	Units (declared)	Meaning
`Molecular_MHCII_kout`	1.9785×10^{-3}	2.3742×10^{-3}	per 5 s	Endosome to surface transport
`Molecular_MHCII_kdegIfng`	4.0846×10^{-5}	4.9015×10^{-5}	per 5 s	IFN-γ decay; unreachable formulation only

Read directly:

Identifier	Value	Units	Meaning
`Molecular_MHCII_tgamInMN1ToMolarConcIFNg`	2.0425×10^{-7}	M per cell (declared)	f_1, applied per Tgam-holding compartment at distance 1. Module
`Molecular_MHCII_tgamInMN2ToMolarConcIFNg`	5.7445×10^{-8}	M per cell (declared)	f_2, applied per Tgam-holding compartment at distance 2. Module
`Molecular_MHCII_mtbToMolarConcAntigen`	4.4917×10^{-8}	M per bacterium	f_a. Module
`Molecular_MHCII_tcellActivationRadius`	1	compartments	Activation radius: the nine-compartment periodic Moore neighbourhood. Module
`Molecular_MHCII_mhcIIBindingThreshold`	150	molecules	Activation threshold on E_out. Active (observation, S9.15)
`Molecular_MHCII_mscale`	44.893	—	σ_m. Module
`Molecular_MHCII_mhcScale`	4.1748	—	σ_H. Module
`Molecular_MHCII_volMIIC`	4.0×10^{-16}	L	Endosomal volume. Module
`Molecular_MHCII_exPepOut`	0	M	External preformed peptide; zero, so no background supply. Module
`Molecular_MHCII_recPerCell`	10^4	molecules	Receptor count; denominator of ω. Read
`Molecular_MHCII_propMhcIn`	0.333333333	fraction	Endosomal share; derivation input. Read
`Molecular_MHCII_propMhcBound`	0.8	fraction	Loaded share; derivation input. Read
`Molecular_MHCII_mhcPerCell`	2.0×10^5	molecules	Total MHC class II; derivation input. Read
`Molecular_MHCII_mrnaPerCell`	10^5	molecules	MHC class II mRNA; derivation input. Read
`Molecular_MHCII_m`, `p`	1.0	ratio	CIITA mRNA and protein initial values; derivation inputs. Read
`Molecular_MHCII_mhcMtbDownReg1` to `4`	0	switch	Down-regulation switches, [0, 1]. Module
`Molecular_MHCII_mtbDownRegRate1` to `4`	0	fraction	r_i. Module
`Molecular_MHCII_mtbHalfDownReg1` to `4`	0	bacteria	K_i. Module
`Molecular_MHCII_ifng`, `exAgOut`	10^{-9}, 10^{-4}	M	Read at construction, overwritten by the first imposition
`Molecular_MHCII_numCells`, `rxnVol`	10^5, 0.001	—	Unreachable formulation only. Inert

The eleven MHC class II initial conditions are read at the construction of every agent; seven of them are derived (M3.6).

Cross-parameter conditions

[1] The fine-grained ceilings must exceed their thresholds (M3.7): `Molecular_TNF_saturationApoptosisTNF` > `Molecular_TNF_thresholdApoptosisTNF`; $\psi_{\text{sat}}(B + R_{\text{IL-10}})$ > `Molecular_IL10_Mac_thresholdSTAT3IL10` for every agent; and `Molecular_TNF_Mac_saturationFracNFkBTNF` $(R_1 + C_1)$ > `Molecular_TNF_Mac_thresholdNFkBTNF`. Not checked.

[2] $\omega \leq 1$ holds provided `Molecular_MHCII_recPerCell` is consistent with the initial receptor condition (M3.6).

[3] If a down-regulation switch is set, its half-saturation must be positive and its rate at most 1 (M3.6). Not checked.

Fixed structural constants

- $\rho_{\text{cell}} = 1.25 \times 10^{11}$ cells L^{-1}, with $V = 8 \times 10^{-12}$ L, so $\phi = 1$ (M3.3).
- κ, compiled, 4.58×10^{-4} s^{-1} (M3.5).
- K_V = 5, the cytoplasm-to-nucleus volume ratio (M3.5).
- The factor 2 on the CXCL9 export (M3.5).
- The 6 s reference of the equilibration step (M3.5).
- The 0.64872 factor relating the initial-pool arguments (M3.5).
- The 5 s reference of the MHC class II rescaling (M3.6).

Read but without effect by default

The neutrophil subgroups `Molecular_TNF_Neutrophil` and `Molecular_IL10_Neutrophil` are read by no code (M9.10.2). The undeclared attribute `TranscriptionFraction` on the `Molecular.TNF` element of the parameter file is never read (Appendix B).

M3.10 Interfaces

With	Provides or receives	See
Spine 1.1	No nested ODD is needed because no per-agent ODE is solved by default	Spine 1.1
Spine 2.1, 2.3	Field units; immediate cross-agent state writes	Spine 2.1, 2.3
Spine Element 3	Molecular substep within each split half-step; list-order processing within a substep	Spine Element 3
Spine 4.1	IFN-γ and antigen are not fields	Spine 4.1
Spine module-slot table	The M3 slot row	Spine 7.14
S1	No change to the movement rule; the NF-κB chemokine output is added in a literal 1 : 1 : 2, which matches the default ratios and so preserves the proportionality movement needs	S1.3, S6.14, M3.5
S2	Macrophage dispatch: MHC class II activation before the suppression test; the NF-κB activation and apoptosis decisions read M3.7 quantities	S2.3, S2.4, M3.6, M3.7
S3	T cell T0 $\rightarrow$ active transition; Treg producer test switch; proliferation still requires an exposed macrophage	S3.5, S3.10, M3.6, M3.7
S4	Live and dead extracellular bacteria feed the MHC class II antigen input	S4, M3.6
S5	Caseum secretes chemokines at a quarter of the macrophage rates, by its own route	S5.6, M3.5
S6	The coarse-grained option: occupancy proxies, decision branches, sinks, secretion helpers and state factors, receptor draws, the derived-chemokine relation, the exclusion of other cell types from the sink	S6.1–S6.14
S8	Recruited T cells start in T0 when MHC class II dynamics is selected	S8.12
S9	MHC class II statistics computed every timestep, and the region centre of mass	S9.15, S9.7
M1	Δt_m and N_m; the rescaling factor; diffusion and dissociation of shed TNFR2; CCL5 and CXCL9 reassignment	M1.2; M1.3 ("Derived assignment of CCL5 and CXCL9", "Post-solve binding sweep")
M2	Recruited T cells start in T0 under either recruitment option when MHC class II dynamics is selected	M2.9, M3.6
M4	The TGF-β activation sweep reads the signalling quantity; the TNF state factor is scaled by bound TGF-β; fibroblasts are integrated by a separate routine in the same loop	M4
M5	Under the three-dimensional option the MHC class II input neighbourhoods are 27, 98 and 27 compartments (9, 16 and 9 in two dimensions), and the conversion factors are not rescaled	M5
M6	Antibiotic killing lowers the intracellular count read by down-regulation	M6
M7	Anti-TNF blocks extend this layout and require the TNFR block; they are unwritten under NF-κB dynamics	M7.8
M8	TNF depletion at the transcription logistic and NF-κB translation; IL-10 depletion at secretion, including the double-secretion path; sinks exempt; Treg induction is a separate slot	M8
M9	Neutrophils are outside the molecular loop and the sink	M9.7
Appendix B	Indexes the four component options, `--odesolver` and the inert parameters	Appendix B