S8. Recruitment

S8.1 Scope and schedule

S8 specifies how new cells enter the grid at the vascular sources under the probabilistic option of M2. Every source is evaluated once per cellular timestep. For each cell type, local soluble signals give an eligibility test and a recruitment weight. A cascade then selects at most one type for each open compartment slot, and the selected cell is constructed and placed.

Probabilistic recruitment from local signals (selected): at each vascular source, each cell type's recruitment probability is a sampled maximum times a fixed weighting of local TNF and chemokine above sampled thresholds. T cell recruitment is gated by an imposed onset and ramp, and a floor restores the macrophage population.

The placement of the sources is specified in spine 5.2, and their properties in spine 2.3. Sources are immovable, carry no state, and are neither created nor destroyed after initialization, so the source list has fixed length and fixed membership for the whole run.

Schedule

Block	Activity	Interval
Initialization, before any other step	Set the T cell onset timestep (S8.10)	Once
Recruitment block: after all movement, before the per-agent statistics reset and the next-state pass (spine 3.2)	Shuffle the source list (S8.2), then make the per-slot attempts at every source (S8.3)	Every cellular timestep
After the statistics block, near the end of the timestep	Latch the T cell onset flag (S8.10)	Every cellular timestep

A recruit has already missed its arrival timestep's movement, so its first movement attempt is at t + 1. The fields recruitment reads are those left by the same timestep's diffusion and molecular block, which has applied that timestep's secretion, transport and decay.

A recruit computes a next state on its arrival timestep, but neither moves nor secretes until the next one. Construction appends it to its type's main agent list (S8.12), and the next-state pass that follows iterates those lists without distinguishing agents created during the timestep. A recruit may therefore kill, suppress, take up bacteria or die on its arrival timestep. Daughter cells, by contrast, are held in separate proliferation lists and do not act on their creation timestep (spine 3.3).

Relationship to oxygen delivery

The same sources deliver oxygen by the flux term at the source compartment (S7.5). Oxygen is consumed by macrophages at their own compartments, wherever they are, not at sources as such (S7.6). Both terms are advanced on the diffusion timestep. Recruitment and oxygen delivery share only the source position list and the caseation test.

Recruitment has no cell-density gate. A source whose oxygen delivery has stopped, at a cell density of 0.8 or more (S7.4), still recruits at full probability, subject only to caseation and a free slot. Neither option of M2 applies such a gate (M2.12). The functionally avascular interior of a lesion (spine 4.2) therefore still admits recruited cells.

Reads and writes

Recruitment reads:

- the source position list, in its current order;
- TNF and CCL2 at the source compartment, and CCL5 and CXCL9 as derived from CCL2 (S6.14);
- the caseation status of the source compartment, its occupant count and maximum occupant count, and whether it holds a large agent or a macrophage;
- the whole-grid macrophage count, for the floor (S8.7);
- the current timestep, the onset timestep and the latched onset flag;
- whether the MHC class II component of M3 is enabled, for the initial T cell state (S8.12);
- the neighbourhood quantities of the neutrophil signal (M9.5);
- the recruitment parameters (S8.13).

Recruitment writes:

- the order of the source list;
- new agents, placed in the first free slot through the admission tests of S8.6 and S8.9, and the per-type agent lists;
- the latched onset flag, in the late call;
- the per-compartment recruitment counters and the recruitment statistics (S8.1, Observation);
- the random number stream (S8.1, Random number consumption).

Recruitment never writes TNF or the chemokines. It affects the fields only through the secretion of the cells it admits, from the next timestep.

Random number consumption

The number and order of draws are part of the model (spine 3.4). The rules that fix the stream are stated at their mechanisms: the per-timestep shuffle (S8.2); the cascade draw taken before any predicate and the second draw for T cells (S8.11); the T cell predicates running in full before onset and the unconditional Tgam draw (S8.8); the override as an assignment (S8.9); the neutrophil crowding draw at every open slot (S8.11, The neutrophil category); and the construction draws (S8.12). The significant-figure truncation applies to threshold comparisons only; probability comparisons use raw values.

Observation

Recruitment writes per-compartment cumulative counters and a group of recruitment statistics: eligible-source, source-active, blocked and caseated-source counts, recruited counts with their cognate partition, and recruitment rates. S9 specifies their definitions, reset codes and output, including the extra-granuloma recruitment statistics (S9.17). The neutrophil recruitment statistics are listed in M9.9.

S8.2 The source list and the order of evaluation

Shuffling

Immediately before the source loop, the source vector is permuted uniformly at random in place. This happens every cellular timestep, six times as often as the agent-list shuffle (`T_shuffle`, spine 3.3). The permutation changes only the order of evaluation; no source moves, and the set of source compartments is invariant.

No two sources share a compartment or a slot, so sources do not compete for space. A recruit at one source also leaves the quantities read at other sources effectively unchanged: it carries no bacteria and secretes nothing until the next timestep's diffusion block.

The order matters through the macrophage floor alone. While the floor is active, each macrophage recruited raises the count the floor reads (S8.7). Sources evaluated early are recruited under the floor, and sources evaluated late revert to the signal-dependent test once the deficit is made good. The shuffle randomizes which sources fall in each group. It also consumes draws, shifting every later draw in the timestep.

The source loop

```
for each source p, in the shuffled order:
    if p is caseated:
        increment the caseated-source statistic
    if n(p) equals the maximum occupant count, or p is caseated:
        skip p
    else:
        make the per-slot attempt (S8.3, S8.11) once per open slot
after the loop: recompute the recruitment rates (S8.1, Observation)
```

The caseated-source statistic is incremented unconditionally, before the skip, so it counts caseated sources, not recruitment failures.

No cell of any type is ever recruited into a caseated compartment. A reimplementation may test caseation once, anywhere, provided this holds.

S8.3 Open slots and the per-slot attempt

Slot accounting

The number of open slots at source p is

$$n_{\mathrm{open}}(p) = N_{\mathrm{slots}} - n(p),$$

with $N_{\mathrm{slots}} = 2$ and $n(p)$ the occupant count of spine 2.3, which counts immune cells and excludes bacteria. The per-slot attempt is made $n_{\mathrm{open}}(p)$ times, with the loop index passed as the slot number: twice at an empty source, once at a source holding one agent, and not at all at a full source.

A later attempt sees the result of an earlier one

The slot count is computed once and is not recomputed. A successful attempt, however, constructs the agent, appends it to its list and places it in the first free slot at once, and every predicate re-reads the occupant count and occupant types. A second attempt therefore sees the result of the first.

Starting from an empty source:

Recruited at the first attempt	At the second attempt	Admissible
Nothing	Any type	Yes, on the same terms as the first
Macrophage	Macrophage	No: a macrophage is a large agent, so the admission predicate fails
Macrophage	T cell	No: the occupancy override (S8.9) assigns false when the single occupant is a macrophage
T cell	T cell	Yes, subject to a fresh draw against `Core_Tcell_probMoveToTcell`
T cell	Macrophage	Yes: a T cell is not a large agent, so the macrophage admission predicate still succeeds

Two cells can therefore be recruited to one compartment in one timestep only as T cell then T cell, or T cell then macrophage. A macrophage recruited at slot 0 closes the compartment to everything.

Use of the slot number

The slot number has one use: each possibility predicate increments its eligible-source statistic only at slot 0, so that statistic counts sources, not slots (S8.1, Observation). The slot number enters no eligibility test, weight, probability or draw. Both attempts use the same parameters and the same functional form, and differ only in the occupancy they observe.

S8.4 Eligibility test and recruitment weight

Four threshold functions, one each for macrophages, Tgam, Tcyt and Treg, share one structure. They differ only in the species they read and in fixed weights. The neutrophil function is built differently (M9.5).

Each function takes the source position and returns an eligibility boolean and a scalar weight. Both are always produced; the calling predicate discards the weight when eligibility is false.

Per-species threshold test

For each species m that the type reads,

$$b_m \iff C_m(p) \geq \theta_m,$$

compared through the significant-figure comparison of spine 2.1. The thresholds come from two sampled parameters per type, one for TNF and one shared by the chemokines:

$$\theta_{\mathrm{TNF}} = \texttt{thresholdRecTNF}, \qquad \theta_{\mathrm{CCL2}} = \texttt{thresholdRecChemokine}, \qquad \theta_{\mathrm{CCL5}} = r_5\,\texttt{thresholdRecChemokine}, \qquad \theta_{\mathrm{CXCL9}} = r_9\,\texttt{thresholdRecChemokine},$$

where each parameter is the type's own (S8.13), and r_5 = `Core_Mac_dCCL5` / `Core_Mac_dCCL2` and r_9 = `Core_Mac_dCXCL9` / `Core_Mac_dCCL2`, the derived-field ratios (S6.14), recomputed inside each function. Each type has one chemokine threshold parameter, scaled per species, not three.

Per-species term

$$w_m = \begin{cases} 0 & \text{if } b_m \text{ is false} \\ \dfrac{V_m\,(C_m(p) - \theta_m)}{(C_m(p) - \theta_m) + K_m} & \text{if } b_m \text{ is true} \end{cases}$$

This is the saturating form H_+ of spine 2.1 applied to the threshold-subtracted concentration. The term is therefore zero at the threshold, and K_m is the excess over threshold at which the term reaches half its supremum. The half-saturation constants are derived like the thresholds:

$$K_{\mathrm{TNF}} = \texttt{recruitmentHalfSatTNF}, \qquad K_{\mathrm{CCL2}} = \texttt{recruitmentHalfSatChemokine}, \qquad K_{\mathrm{CCL5}} = r_5\,\texttt{recruitmentHalfSatChemokine}, \qquad K_{\mathrm{CXCL9}} = r_9\,\texttt{recruitmentHalf}$$

Eligibility and weight

$$\text{eligible} \iff \bigvee_m b_m, \qquad W = \sum_m w_m.$$

One species above its threshold suffices for eligibility; every other species then contributes zero. Each w_m is at most V_m, and the V_m sum to 1 for every type, so W is at most 1, reaching 1 only in the limit where every species read is far above its threshold. W is non-negative except near a threshold, where truncation can admit a term below zero of order 10^{-7} relative.

The fixed weights

The V_m are literal constants, not parameters. TNF receives a fixed share f_{TNF} of the unit weight, and the remainder is divided equally among the chemokines the type reads.

Type	CCL2	CCL5	CXCL9	TNF (f_{TNF})
Macrophage	3/8	3/8	not read	1/4
Tgam	2/9	2/9	2/9	1/3
Tcyt	not read	1/3	1/3	1/3
Treg	not read	1/4	not read	3/4

The TNF share is an independent choice per type, made in the source and not sampled. It alone sets the balance between TNF-driven and chemokine-driven recruitment for each type.

How the weight is used

The calling predicate forms a probability and compares it with a fresh uniform draw:

$$P_{\mathrm{mac}} = \texttt{maxRecProb} \cdot W, \qquad P_{\mathrm{Tcell}} = \texttt{maxRecProb} \cdot W \cdot f(t),$$

with `maxRecProb` the sampled maximum of the type and $f(t)$ the ramp factor of S8.10. Success makes the type possible. It does not recruit the cell, because the cascade (S8.11) may then select another type.

S8.5 Collapse of the chemokine terms

Under the derived-field scheme every chemokine term of S8.4 reduces to the same function of CCL2, so all types share one chemokine sensitivity profile.

After each solve, CCL5 and CXCL9 are set pointwise as $C_5 = r_5 C_2$ and $C_9 = r_9 C_2$ (S6.14), with the same ratios that scale the thresholds and half-saturations of S8.4. Writing θ_c and K_c for the unscaled chemokine threshold and half-saturation,

$$w_{\mathrm{CCL5}} = \frac{V_5\,r_5\,(C_2 - \theta_c)}{r_5\,[(C_2 - \theta_c) + K_c]} = \frac{V_5\,(C_2 - \theta_c)}{(C_2 - \theta_c) + K_c},$$

and likewise for CXCL9 with r_9. Each chemokine term is the CCL2 term times that chemokine's weight. For $r > 0$ the three chemokine booleans also coincide in exact arithmetic, since $r_5 C_2 \geq r_5 \theta_c \iff C_2 \geq \theta_c$. Under the significant-figure comparison the collapse is exact for CCL5, because $r_5 = 1$ makes the CCL5 field identical to CCL2. For CXCL9 it holds to within a relative tolerance of order 10^{-7}, near a threshold only.

Collapsed form

$$W = (1 - f_{\mathrm{TNF}}) \left[\frac{C_2 - \theta_c}{(C_2 - \theta_c) + K_c} \right]_+ + f_{\mathrm{TNF}} \left[\frac{C_{\mathrm{TNF}} - \theta_T}{(C_{\mathrm{TNF}} - \theta_T) + K_T} \right]_+ ,$$

where $[\cdot]_+$ is the quantity itself when the threshold is met and zero otherwise.

Type	Chemokine coefficient	TNF coefficient
Macrophage	3/4	1/4
Tgam	2/3	1/3
Tcyt	2/3	1/3
Treg	1/4	3/4

Each type's weight is therefore a function of CCL2 and TNF at the source, with four sampled quantities (two thresholds, two half-saturations) and one fixed split.

Consequences

The chemokine coefficient is $1 - f_{\mathrm{TNF}}$ whatever the number of chemokines a type reads. A reimplementation may replace the CCL5 and CXCL9 reads with one CCL2 term of coefficient $1 - f_{\mathrm{TNF}}$. It then reproduces this model exactly for CCL5, and for CXCL9 except within the near-threshold tolerance. Chemotaxis differs: there the per-type sensitivities still scale the magnitude of the response (S1.3), whereas in recruitment the scaling cancels.

The collapse is contingent. It holds because the ratios that scale the thresholds are the ratios that derive the fields, both computed from the same three macrophage secretion-rate parameters. It would fail if the chemokines had independent diffusivities or degradation rates, if a secretor other than the macrophage produced them in another proportion, or if the M3 NF-κB chemokine output wrote them in another ratio. The first two arise in no configuration; for the third, see M3. The three secretion rates are read on every call only to form the ratios, so they have no effect on any weight; they are fixed by default, so r_5 and r_9 are constant across the virtual library.

S8.6 Macrophage eligibility and recruitment probability

The macrophage threshold function reads CCL2, CCL5 and TNF, not CXCL9 (S8.4). Its four parameters are `Recruit_Prob_Mac_thresholdRecChemokine`, `Recruit_Prob_Mac_thresholdRecTNF`, `Recruit_Prob_Mac_recruitmentHalfSatChemokine` and `Recruit_Prob_Mac_recruitmentHalfSatTNF`, all sampled.

The predicate

```
macrophage predicate at source p, slot s:
    (eligible, W) <- macrophage threshold function at p       // always evaluated
    if eligible and s == 0:
        increment the macrophage eligible-source statistic
    blocked <- n(p) == N_slots or not macrophage admission at p
    if blocked:
        increment the macrophage blocked-source statistic
    bigBlocked <- n(p) < N_slots and a large agent is at p
    if bigBlocked:
        increment the macrophage blocked-by-large-agent statistic
    if p is caseated or blocked:
        return false
    if whole-grid macrophage count < Core_Mac_initNumber:    // the floor, S8.7
        return true                                           // no draw
    return eligible and U(0,1) < Recruit_Prob_Mac_maxRecProb * W
```

The order matters: the statistics are written before the early returns, and the probability draw is reached on one branch only.

Macrophage admission at recruitment is the macrophage movement admission predicate itself (S1.4), so a reimplementation should share the predicate. It requires that no large agent be present. Macrophages are the only large agents by default, so a source holding a macrophage never recruits another, although it has a free slot. T cell recruitment, by contrast, is stricter than T cell movement (S8.9).

When the floor fires, the predicate returns true without consulting the eligibility boolean and without a draw. The threshold function has already run, and its statistic has been written.

S8.7 The macrophage recruitment floor

The rule

The floor reads the whole-grid macrophage count statistic. That count includes every state, the dead state included, and recruits placed earlier in the same block. While it is strictly below `Core_Mac_initNumber`, the macrophage is a candidate unconditionally, with no signal test and no draw. The eligibility boolean is not consulted and the weight is unused. Only the conditions tested before the floor still apply: the compartment must not be caseated and must not be blocked, that is, it must have a free slot and no large agent.

The floor is therefore a bypass of the whole signal mechanism, not a change to a probability. While it is active, local TNF and chemokines are irrelevant to macrophage recruitment.

How the count evolves within a timestep

Construction calls the per-agent statistics update on the recruit, registering it in the macrophage count at once, and the agent-statistics reset runs after the recruitment block. Within one block the count therefore holds the value committed at the end of the previous timestep plus every macrophage recruited so far. Because statistics are written before dead agents are erased, macrophages that died at the previous commitment count toward the floor.

The floor therefore closes progressively as sources are visited. The population does not overshoot it: once enough sources have fired to restore the count, the remaining sources revert to the signal test.

With the floor active, macrophage possibility holds at every non-caseated source with a free slot and no macrophage, which outside the lesion is very nearly all 7886 sources. Before T cell onset no other category can be possible at a slot (the neutrophil category never is; S8.11, The neutrophil category), so the cascade's share is 1 and recruitment at an admissible source is certain. A deficit is normally made good within the timestep in which it is first observed.

At the opening timestep the count holds the initial resting population plus the seeding infected macrophage (spine 5.4). It therefore exceeds the floor by one, and the floor first fires once the population has fallen by two or more. This holds because `Core_Mac_initNumber` counts only the resting macrophages placed by density.

The floor value

`Core_Mac_initNumber` is computed at load from `Core_Mac_initDensity` and the grid size, and a value supplied in the parameter file is overwritten (spine 5.3). With the density sampled over [0.004799, 0.010262], the floor lies between 431 and 923. It varies across the virtual library in lockstep with the initial population, not independently of it.

The floor applies under both options of M2 (M2.9). It is an imposed quantity, not an emergent one; its interpretation is specified in spine 4.2.

S8.8 T cell eligibility and recruitment probability

The three T cell threshold functions read the species of S8.4: Tgam reads CCL2, CCL5, CXCL9 and TNF; Tcyt reads CCL5, CXCL9 and TNF; Treg reads CCL5 and TNF. By the collapse of S8.5, Tgam and Tcyt are identical in form and differ only in their parameter values, while Treg carries three quarters of its weight on TNF. Each type takes its four parameters from its own group, `Recruit_Prob_Tcell_<type>`; all twelve are sampled.

The predicate

```
T cell predicate for type X at source p, slot s, timestep t:
    f <- ramp factor at t (S8.10)                    // no draw
    (eligible, W) <- threshold function of X at p
    if eligible and s == 0:
        increment the X eligible-source statistic
    if p is caseated or n(p) == N_slots:
        return false
    if n(p) == 1:                                     // occupancy override, S8.9
        if a macrophage is at p:
            eligible <- false
        else:
            eligible <- U(0,1) < Core_Tcell_probMoveToTcell
    return eligible and U(0,1) < Recruit_Prob_Tcell_X_maxRecProb * W * f
```

The ramp factor is computed at the top of the predicate, at every attempt, before the threshold test and the caseation gate.

The onset flag is not tested in the predicate. The caller conjoins it on the right of a short-circuiting AND (S8.11), so the predicate runs in full, and takes its draws, from the opening timestep. The ramp factor is 0 before onset, so the final draw cannot succeed then.

The Tgam draw

The Tgam predicate takes its probability draw unconditionally, into a variable, before forming the conjunction. The Tcyt and Treg predicates draw only when their boolean, possibly overridden, is true. The returned values agree in distribution, but within the non-caseated, non-full branch Tgam always draws, so the draw streams differ. A reimplementation that regularized the three bodies would diverge from this model.

S8.9 The occupancy override

The rule

At a source holding exactly one agent, the eligibility boolean is replaced:

$$\text{eligible} \leftarrow \begin{cases} \text{false} & \text{if the occupant is a macrophage} \\ U(0,1) < \texttt{Core_Tcell_probMoveToTcell} & \text{otherwise} \end{cases}$$

The replacement is an assignment, not a conjunction. The recruitment weight W is not overwritten and still sets the final probability, so a source with one T cell keeps its signal-dependent probability behind a signal-independent gate.

The occupant test counts macrophages. With exactly one occupant, a count of zero identifies a T cell, the only other class present by default. With neutrophils or fibroblasts present, the second branch applies to them too, so `Core_Tcell_probMoveToTcell` governs recruitment beside them (M9.6, M4).

Consequences

A source holding a macrophage recruits no T cell (S8.3). Sources in the macrophage core, where TNF and chemokine signal is strongest, carry the highest T cell weight and are the most likely to hold a macrophage. T cells reach the core only by movement after recruitment elsewhere.

The outcome equals that of a conjunction: a false threshold boolean implies W = 0, so the override cannot make a sub-threshold source possible. At a source holding one T cell the effective probability is

$$P = \texttt{Core_Tcell_probMoveToTcell} \cdot \texttt{maxRecProb} \cdot W \cdot f(t),$$

so with `Core_Tcell_probMoveToTcell` = 0.08 one T cell occupant cuts the probability by a factor of 12.5. The assignment form nonetheless changes which later draws are taken.

Recruitment is stricter than movement. Movement admits a T cell beside a macrophage with probability `Core_Tcell_probMoveToMac` and beside a T cell with `Core_Tcell_probMoveToTcell` (S1). Recruitment forbids the first and permits the second on the same terms. The reverse pair, a macrophage recruited beside a T cell, is admitted (S8.3). The spine records this among the model's interaction asymmetries (spine 4.8).

S8.10 Onset of T cell recruitment and the ramp

T cell recruitment is gated in three stages, implemented in different places: an onset timestep set at initialization, a flag latched when it is reached, and a ramp factor.

The onset timestep

The onset timestep is set once, at the top of the initialization routine, before any entity is placed:

$$t_{\mathrm{rec}} = \begin{cases} \texttt{Recruit_Prob_timeRecEnabled} & \text{if } \texttt{Recruit_Prob_timeRecEnabledRange} = 0 \\ \lfloor U(\texttt{timeRecEnabled} - \texttt{timeRecEnabledRange},\ \texttt{timeRecEnabled} + \texttt{timeRecEnabledRange}) \rfloor & \text{otherwise} \end{cases}$$

The second branch is a continuous uniform draw truncated toward zero. With `Recruit_Prob_timeRecEnabledRange` = 0 the first branch is taken and no draw is made. A reimplementation with a non-zero range should use the continuous draw with truncation, not a discrete uniform, since the two give different probabilities to the endpoints. With `Recruit_Prob_timeRecEnabled` sampled over [3160, 4275] timesteps, t_{rec} lies between about 22 and 30 days.

The latched flag

The latch is called once per cellular timestep, after the statistics block. It returns at once if the flag is already set, so the flag is monotone: once set, it is never cleared. Exactly one of two criteria is active, chosen by a parameter value:

$$\text{set the flag if } \begin{cases} n_{\mathrm{Mtb}}^{\mathrm{total}} > \texttt{Recruit_Prob_mtbRecEnabled} & \text{if } \texttt{mtbRecEnabled} > 0 \\ t \geq t_{\mathrm{rec}} & \text{otherwise} \end{cases}$$

Here $n_{\mathrm{Mtb}}^{\mathrm{total}}$ is the total extracellular plus total intracellular bacterial count; the intracellular total includes neutrophil-held bacteria (M9.9). With `Recruit_Prob_mtbRecEnabled` = 0 the time criterion is taken. The bacterial branch is inert by default but live: a positive value would make the onset of adaptive immunity a response to the infection rather than an imposed schedule (spine 4.3). Were it selected, it would read the counts from the same timestep's statistics block, after that timestep's bacterial growth and killing.

Because the latch follows the recruitment block, a flag set at timestep t is first read at t + 1. The effective onset is therefore $t_{\mathrm{rec}} + 1$.

The ramp factor

$$f(t) = \begin{cases} 0 & t < t_{\mathrm{rec}} \\ 1 & t \geq t_{\mathrm{rec}} \text{ and } \tau \leq 0 \\ \min\left(\max\left(\dfrac{t - t_{\mathrm{rec}}}{\tau},\ 0\right),\ 1\right) & t \geq t_{\mathrm{rec}} \text{ and } \tau > 0 \end{cases}$$

with τ = `Recruit_Prob_timeMaxRecDelay`. The quotient is computed in floating point on integer operands, so the ramp advances in steps of $1/\tau$. With τ sampled over [770, 1196] timesteps, full rate is reached 5 to 8 days after onset, or 27 to 38 days after infection.

The ramp counts from t_{rec}, not from the latch, so under the time criterion, at the first timestep on which T cell recruitment is permitted, $f = 1/\tau$, not 0. A reimplementation counting from the latch would differ. Under the bacterial criterion the flag could latch before t_{rec}, but f stays 0 until t_{rec} and then ramps from it, so that criterion can delay recruitment beyond t_{rec} but cannot advance it.

Where each stage applies

Stage	Evaluated	Effect when not satisfied
Onset timestep	Once, at initialization	None; it is a stored value
Latched flag	Once per timestep, after the statistics block	Conjoined by the caller with each T cell predicate's result (S8.11), after the predicate has run
Ramp factor	At every slot attempt, at the top of each T cell predicate	f = 0, so the probability draw cannot succeed

Before onset the T cell predicates therefore run in full and take their draws, and the cascade discards their results (S8.8).

Macrophage recruitment has no onset gate and no ramp. It is active from the opening timestep, subject to the thresholds of S8.6 and the floor of S8.7. Only the gating of this section delays T cell recruitment relative to macrophage recruitment.

S8.11 The recruitment cascade

The cascade runs once per open slot. It finds the possible types, divides a unit probability among them, and recruits at most one cell.

The routine

```
cascade at source p, slot s:
    r <- U(0,1)                                  // before any predicate
    pmac   <- macrophage predicate(p, s)
    pneut  <- neutrophil predicate(p, s)
    ptgam  <- Tgam predicate(p, s) and onset flag
    ptcyt  <- Tcyt predicate(p, s) and onset flag
```

```
ptreg  <- Treg predicate(p, s) and onset flag
ptcell <- ptgam or ptcyt or ptreg
add each possibility (0 or 1) to its source-active statistic
k     <- pmac + pneut + ptcell                  // T cells count as one
delta <- 1 / k, or 0 if k == 0
run   <- 0
if pmac:
    run <- run + delta
    if r <  run: recruit a macrophage at p; return
if pneut:
    run <- run + delta
    if r <  run: recruit a neutrophil at p (M9.5); return
if ptcell:
    run <- run + delta
    if r <= run:
        kT <- ptgam + ptcyt + ptreg
        deltaT <- 1 / kT, or 0 if kT == 0
        runT <- 0
        rT <- U(0,1)                            // second, independent draw
        if ptgam: runT <- runT + deltaT; if rT <  runT: recruit a Tgam; return
        if ptcyt: runT <- runT + deltaT; if rT <  runT: recruit a Tcyt; return
        if ptreg: runT <- runT + deltaT; if rT <= runT: recruit a Treg; return
        fatal error
// nothing recruited at this slot
```

Division among possible categories

Dividing the unit probability replaces an enumeration of type combinations, eight for three categories and sixteen for four. The rationale is that the division avoids case analysis. Each type's realized rate depends on how many other types are possible at the same slot; no parameter expresses this dependence.

The three T cell types count as one category. The T cell share $1/k$ is divided evenly among the T cell types that are possible. Given that a type is possible, its chance of selection is:

Possible types	k	Macrophage	Each possible T cell type
Macrophage only	1	1	—
One T cell type only	1	—	1
Macrophage and one T cell type	2	1/2	1/2
Macrophage and two T cell types	2	1/2	1/4
Macrophage and three T cell types	2	1/2	1/6
Three T cell types only	1	—	1/3

A macrophage competing against all three T cell types is three times as likely to be selected as any one of them. This is structural: the per-type `maxRecProb` governs whether a type is possible, not how the probability is then divided. The neutrophil category is counted separately; it is never possible by default (S8.11, The neutrophil category), so the divisions are exactly those tabulated.

Comparisons and outcomes

The macrophage, neutrophil, Tgam and Tcyt branches compare strictly. The outer T cell branch and the Treg branch, the final branch of each cascade, compare non-strictly, as a catch-all for a draw equal to the running total. With a draw on [0, 1) the catch-all never binds and the fatal error is unreachable.

If no type is possible, $k = 0$, the share is 0, no branch is entered and nothing is recruited. This is the only way a slot attempt recruits nothing. Because the draw is on [0, 1), if any type is possible exactly one cell is recruited. The chance that a slot recruits nothing therefore lies wholly in the predicates. The realized probability that a type is recruited is the probability that it is possible times its cascade share, which depends on which other predicates succeeded at that slot.

The neutrophil category (M9)

The neutrophil predicate is the cascade's second category. It is evaluated at every open slot of every source that the loop does not skip, for the whole run, and takes one crowding draw at each, in every configuration.

The predicate, its signal, its depletion window and the creation of a recruited neutrophil are specified in M9.5. In the default configuration the window refuses neutrophil recruitment at every timestep, so the category is never possible (M9.1, M9.5).

S8.12 Construction of a recruited cell

Five routines construct a recruited cell, one per type. They share a structure: check the preconditions, increment the per-compartment counters, construct the agent, and update the statistics. The simulation-level factory functions append the agent to its type's list and place it through the admission predicate of its type (S8.6, S8.9). Placement re-applies the predicate and stops the run with a fatal error if there is no room, so no agent is ever placed in violation of the occupancy rules.

Two per-compartment counters are incremented at the source, a general recruitment counter and a per-type counter. They are cumulative over the run, never reset, and read by no rule (S9).

Birth time and age at arrival

A recruited cell is given a birth time in the past, so it arrives with a random age and a correspondingly reduced residual lifespan.

Type	Birth time	Age at arrival	Mean residual lifespan
Macrophage	$t - U\{0, \dots, \texttt{Core_Mac_maxAge} - 1\}$	0 to 14399	≈ 7200 timesteps, 50 d
Tgam, Tcyt, Treg	$t - U\{1, \dots, \texttt{Core_Tcell_maxAge} - 1\}$	1 to 431	≈ 216 timesteps, 1.5 d

The integer draw is half-open at its upper bound, so no recruit arrives at its maximum age, and a recruited macrophage can arrive at age 0. Senescence is tested as $t \geq$ death time, with death time equal to birth time plus maximum age, so every recruit survives at least to the next timestep.

A recruit created at one below its maximum age computes a next state on its arrival timestep (S8.1). At the next timestep it has its last movement opportunity (a T cell always attempts; a macrophage only if its age passes the interval test), and it senesces in that timestep's next-state computation.

Because the mean age at arrival is half the maximum, a recruitment event contributes on average about half the maximum lifespan in cell-time. The randomized birth time also desynchronizes the age-based movement attempts of recruited macrophages (S1.2); a birth time of t for every recruit would synchronize them.

Antigen specificity

Each T cell construction draws once against its type's specificity probability:

$$\text{cognate} \iff U(0,1) < \texttt{Recruit_Prob_Tcell_<type>_probCognate}.$$

The draw is taken during construction, after the cascade has selected the type, so it does not affect which type is recruited. Specificity is fixed at construction and never revised; its consequences are specified in S3.2. With the sampled ranges, between about one in 190 and one in ten recruited T cells is cognate. Macrophages and neutrophils take no specificity draw.

Initial state

Type	Initial state	With the MHC class II component of M3 enabled
Macrophage	Resting, no intracellular bacteria, both pro-inflammatory flags clear	Unchanged
Tgam	Active	T0
Tcyt	Active	T0
Treg	Active	T0

With the MHC class II component disabled, as by default, T cells arrive active (M3.6). This is the only module dependence that changes the state in which a recruit is constructed, and the reason the T0 states are declared but unreachable in the default configuration (spine 2.5).

A recruited macrophage is constructed with the simulation's growth-rate randomization and NF-κB dynamics settings, which govern the draws its constructor makes (S6.3, M8).

S8.13 Parameters

Identifier	Value or range	Range	Units	Meaning
`Recruit_Prob_Mac_maxRecProb`	0.0881 to 0.3083	Sampled	probability	Maximum macrophage recruitment probability (S8.4, S8.6)
`Recruit_Prob_Mac_thresholdRecChemokine`	0.9054 to 3.1098	Sampled	molecules per compartment[1]	Chemokine threshold and subtrahend (S8.4)
`Recruit_Prob_Mac_thresholdRecTNF`	0.007491 to 0.062595	Sampled	molecules per compartment[1]	TNF threshold and subtrahend (S8.4)
`Recruit_Prob_Mac_recruitmentHalfSatChemokine`	2.3725 to 4.6218	Sampled	molecules per compartment[1]	Chemokine half-saturation excess (S8.4)
`Recruit_Prob_Mac_recruitmentHalfSatTNF`	1.8389 to 5.0538	Sampled	molecules per compartment[1]	TNF half-saturation excess (S8.4)
`Recruit_Prob_Tcell_Tgam_maxRecProb`	0.1565 to 0.8374	Sampled	probability	Maximum Tgam recruitment probability (S8.8)
`Recruit_Prob_Tcell_Tgam_thresholdRecChemokine`	0.03667 to 0.09350	Sampled	molecules per compartment[1]	Chemokine threshold (S8.4)
`Recruit_Prob_Tcell_Tgam_thresholdRecTNF`	0.8878 to 1.2104	Sampled	molecules per compartment[1]	TNF threshold (S8.4)
`Recruit_Prob_Tcell_Tgam_recruitmentHalfSatChemokine`	0.9830 to 1.5233	Sampled	molecules per compartment[1]	Chemokine half-saturation excess (S8.4)
`Recruit_Prob_Tcell_Tgam_recruitmentHalfSatTNF`	0.7497 to 1.6375	Sampled	molecules per compartment[1]	TNF half-saturation excess (S8.4)
`Recruit_Prob_Tcell_Tgam_probCognate`	0.02645 to 0.09421	Sampled	probability	Specificity draw for a recruited Tgam (S8.12)
`Recruit_Prob_Tcell_Tcyt_maxRecProb`	0.1815 to 0.8188	Sampled	probability	Maximum Tcyt recruitment probability (S8.8)
`Recruit_Prob_Tcell_Tcyt_thresholdRecChemokine`	2.9318 to 4.0693	Sampled	molecules per compartment[1]	Chemokine threshold (S8.4)
`Recruit_Prob_Tcell_Tcyt_thresholdRecTNF`	0.9761 to 1.5167	Sampled	molecules per compartment[1]	TNF threshold (S8.4)
`Recruit_Prob_Tcell_Tcyt_recruitmentHalfSatChemokine`	4.7572 to 10.2728	Sampled	molecules per compartment[1]	Chemokine half-saturation excess (S8.4)
`Recruit_Prob_Tcell_Tcyt_recruitmentHalfSatTNF`	0.9764 to 1.5277	Sampled	molecules per compartment[1]	TNF half-saturation excess (S8.4)
`Recruit_Prob_Tcell_Tcyt_probCognate`	0.005940 to 0.105197	Sampled	probability	Specificity draw for a recruited Tcyt (S8.12)
`Recruit_Prob_Tcell_Treg_maxRecProb`	0.1838 to 0.8140	Sampled	probability	Maximum Treg recruitment probability (S8.8)
`Recruit_Prob_Tcell_Treg_thresholdRecChemokine`	1.4825 to 2.5441	Sampled	molecules per compartment[1]	Chemokine threshold (S8.4)
`Recruit_Prob_Tcell_Treg_thresholdRecTNF`	0.9445 to 2.0493	Sampled	molecules per compartment[1]	TNF threshold (S8.4)
`Recruit_Prob_Tcell_Treg_recruitmentHalfSatChemokine`	0.9403 to 2.0387	Sampled	molecules per compartment[1]	Chemokine half-saturation excess (S8.4)
`Recruit_Prob_Tcell_Treg_recruitmentHalfSatTNF`	1.4596 to 2.3565	Sampled	molecules per compartment[1]	TNF half-saturation excess (S8.4)

Identifier	Value or range	Range	Units	Meaning
`Recruit_Prob_Tcell_Treg_probCognate`	0.005366 to 0.104857	Sampled	probability	Specificity draw for a recruited Treg (S8.12)
`Recruit_Prob_timeRecEnabled`	3160 to 4275	Sampled	integer, timesteps (22–30 d)	Onset timestep of T cell recruitment (S8.10)[2]
`Recruit_Prob_timeMaxRecDelay`	770 to 1196	Sampled	timesteps (5–8 d)	Ramp duration (S8.10)
`Recruit_Prob_mtbRecEnabled`	0	—	bacteria	Bacterial onset threshold; 0 selects the time criterion (S8.10)
`Recruit_Prob_timeRecEnabledRange`	0	—	timesteps	Half-width of the onset draw; 0 means no draw (S8.10)

[1] Declared without units and compared with field values, so implicitly molecules per compartment (spine 2.1). The chemokine parameters are CCL2-equivalent values, scaled per species (S8.4). These four parameters per type have no declared range and are not checked; `maxRecProb` and `probCognate` are range-checked to [0, 1].

[2] Declared with a minimum of 0 and no maximum, so a value beyond the run would suppress T cell recruitment throughout; no sampled value does.

No cross-parameter consistency constraint involves an S8 parameter.

Fixed structural constants: the TNF shares 1/4 (macrophage), 1/3 (Tgam, Tcyt) and 3/4 (Treg), with the remainder divided equally among the chemokines a type reads (S8.4); the override trigger, an occupant count of exactly 1, and its value `false` at a macrophage-occupied source (S8.9); the statistics gate, slot number 0 (S8.3); the cascade division 1/k among possible categories, with T cells as one category (S8.11); strict comparisons except at the final branch of each cascade (S8.11); $N_{\text{slots}} = 2$ (spine 2.3).

Owned elsewhere and read at recruitment: `Core_Tcell_probMoveToTcell` sets the override draw (S8.9; owned by S1); `Core_Mac_maxAge` and `Core_Tcell_maxAge` bound the birth-time draws (S8.12; owned by S2 and S3).

Read here but without effect on recruitment: `Core_Mac_dCCL2`, `Core_Mac_dCCL5` and `Core_Mac_dCXCL9`, read on every call to form ratios that cancel (S8.5; S6 owns them); the neutrophil recruitment parameters, read at every slot attempt with nil effect by default (M9.10.2). The `Recruit_ODE` and `Recruit_OLD` groups are not read under this option; option B reads `Recruit_ODE` and `Recruit_OLD_Mac_probRec` (M2.9; Appendix B).

S8.14 Interfaces

With	Provides or receives	See
Spine 2.1	Receives the field unit, the saturating form H_+ and the significant-figure comparison	S8.4
Spine 2.3	Receives the occupancy rules, slot count, size classes and source properties	S8.1, S8.3, S8.6
Spine 2.5	Receives the T cell state inventory, including the T0 states unreachable by default	S8.12
Spine 3.2, 3.3	Receives the position of the recruitment block and the daughter deferral	S8.1
Spine 4.2	Holds the interpretation of the macrophage floor as an imposed quantity	S8.7
Spine 4.3, 4.6, 4.8	Records the onset design (4.3), the whole-grid read of the floor (4.6) and the recruitment–movement asymmetry (4.8)	S8.7, S8.9, S8.10
Spine 5.2, 5.4	Receives the source list and the seeding macrophage	S8.1, S8.7
Spine 5.2, 5.3	Receives `Core_nrSources` (5.2) and `Core_Mac_initNumber` (5.3)	S8.7
S1	Owns `Core_Tcell_probMoveToTcell` and `Core_Tcell_probMoveToMac`, the chemotactic sensitivities and the macrophage movement admission predicate (S1.4); recruits make no movement attempt on arrival	S8.1, S8.5, S8.6, S8.9
S2	Owns `Core_Mac_maxAge` and macrophage senescence; receives recruited macrophages	S8.12
S3	Owns `Core_Tcell_maxAge`, the consequences of specificity and producer status; receives recruited T cells	S8.12
S5	A recruit may be killed by the caseation cascade on its arrival timestep	S8.1
S6	Owns the derived chemokine fields and ratios (S6.14), the macrophage secretion rates and the construction draws (S6.3)	S8.4, S8.5, S8.12
S7	Shares the source list and the caseation test; holds the density gate that recruitment lacks	S8.1
S9	Owns the recruitment statistics, counters and rates	S8.1
M1	The collapse holds only while CCL5 and CXCL9 are set from CCL2 after each solve, as every solver does (M1.3, Derived assignment of CCL5 and CXCL9); independent diffusion of the three chemokines would restore three terms per type	S8.5
M2	S8 specifies the selected option (slot marker in S8.1). The alternative, lymph-node and blood recruitment, replaces S8.4 to S8.12 except the floor of S8.7, the eligibility tests of S8.4, S8.6 and S8.8 (without the weight), the collapse of S8.5 and the T0 rule of S8.12 (M2.9); under it the onset and magnitude of the T cell response are outputs	M2.9
M3	With the MHC class II component enabled, recruited T cells are constructed in T0; an NF-κB chemokine output in another ratio would break the collapse	S8.5, S8.12
M5	The source count follows the source density and the grid size; the neighbourhood size in the neutrophil signal is 9 in two dimensions and 27 in three	S8.1, M9.5

With	Provides or receives	See
M6	Adds drug flux at the same sources on the diffusion timestep; reads or alters no recruitment quantity	S8.1
M7	Adds anti-TNF flux at the same sources; alters recruitment only through the TNF field	S8.4
M8	Neutrophil injection and reinfection create cells outside the cascade and increment no recruitment counter; the reinfecting macrophage counts toward the floor (M8.3); growth-rate randomization governs construction draws	S8.7, S8.12
M9	Owns the neutrophil predicate, signal, depletion window and creation; neutrophil-held bacteria enter the bacterial onset criterion	S8.10, S8.11
Appendix B	Indexes the recruitment toggles and the never-read recruitment parameters	S8.13

Dendritic cells are absent from the default configuration (spine 2.3). The selected option has no dendritic-cell predicate or construction routine, so it can recruit none; the per-type recruitment statistics nonetheless include a dendritic-cell entry, identically zero (S9).