HIDDEN RESERVOIRS: HOW HOST GENETIC HETEROGENEITY DEFEATS SYMPTOM-DRIVEN MALARIA SURVEILLANCE

SÁNDOR BARTHA, AKINDELE A. ONIFADE, DAN SHEARER, AND WILLIAM WAITES

Abstract. Malaria surveillance in endemic settings relies predominantly on symptom-driven testing. In populations carrying sickle cell trait (HbAS), this strategy has a structural blind spot: HbAS carriers experience reduced clinical severity while remaining infectious, making them largely invisible to symptom-triggered diagnostics. We use a rule-based stochastic model implemented in the Kappa language to characterise the interaction between host genetic heterogeneity and surveillance bias. A systematic parameter sweep across reservoir size, test bias, and test rate reveals that the interaction is multiplicative: at realistic HbAS prevalence (20–25%), purely symptomatic testing requires roughly double the effort of random testing to achieve epidemic suppression, and at higher prevalence fails entirely regardless of intensity. Sensitivity analysis confirms that this effect is driven by observation bias — who gets tested — rather than by differential recovery rates. Quantified comparisons across testing strategies show that hybrid approaches allocating a substantial fraction of tests randomly outperform both purely symptomatic and purely random strategies. The test count analysis demonstrates that symptom-biased testing is a false economy: it consumes fewer tests but fails to control transmission, while random testing costs more but succeeds. These results suggest that incorporating random screening into malaria surveillance programmes in HbAS-prevalent regions could substantially improve epidemic control at modest additional cost.

1. Introduction

Malaria remains a major and worsening global health burden, particularly in Sub-Saharan Africa, where *Plasmodium falciparum* is responsible for the most severe cases. Despite advances in insecticide-treated nets, diagnostics, and artemisinin-based therapies, the disease caused an estimated 282 million cases and 610,000 deaths in 2024 — an increase of nine million cases on the previous year [1]. The burden is heavily concentrated: eleven countries account for roughly two thirds of all cases and deaths, with approximately 31% of global malaria mortality in Nigeria alone [2]. Against this backdrop, a continental vaccination campaign is now underway: as of early 2026, 25 African countries have introduced malaria vaccines (RTS,S and R21/Matrix-M) into routine childhood immunisation programmes, with Nigeria beginning its phased rollout in late 2024 [3]. For the four human-adapted species (*P. falciparum*, *P. vivax*, *P. malariae*, and *P. ovale*), humans are the sole vertebrate host. *P. falciparum* has no dormant liver stage and no animal reservoir; it must maintain continuous asexual replication in human blood throughout the dry season (up to six months in the West African Sahel) making the human infectious reservoir the sole mechanism by which the parasite bridges transmission seasons [4, 5]. Unlike zoonotic species such as *P. knowlesi* [6], the human-adapted species can in principle be eliminated by targeting the human reservoir alone. Any human subpopulation harbouring transmissible gametocytes while evading detection therefore represents a blind spot in malaria control. The most epidemiologically consequential such subpopulation is asymptomatic carriers whose prevalence and transmission potential are shaped in part by host genetics [7, 8].

Sickle cell trait (HbAS) creates such a blind spot. HbAS substantially protects against severe malaria without preventing infection [9–13]. Carriers yet remain susceptible to blood-stage infection and may harbour parasites at low or submicroscopic densities without overt symptoms. Lawaly et al. [14] demonstrated a significant genetic contribution to gametocyte prevalence specifically in asymptomatic infections, the population which symptom-driven surveillance misses. Gonçalves et al. [15] confirmed that HbAS and HbAC carriers harbour more gametocytes, sustain longer infections, and are more infectious in direct skin-feeding assays; HbAC individuals comprising 5% of the study population contributed approximately a third of the sexual-stage reservoir at peak transmission. In a Ugandan study, HbAS

Date: March 10, 2026.

individuals were nearly three times more likely to carry detectable gametocytes (HR = 2.68) at densities an order of magnitude higher (DR = 9.19) than HbAA individuals [8]. Ngou et al. [16] confirmed twice the mosquito infection risk from HbAS donors in membrane-feeding assays. Infected HbAS carriers are thus more infectious, and largely asymptomatic.

WHO guidelines define a suspected malaria case as any individual presenting with fever or a history of fever with no other obvious cause. Passive case detection — testing patients who self-present to health facilities with symptoms — is accordingly the primary surveillance mechanism in high-burden countries [17, 18], with only two thirds of febrile episodes leading to treatment-seeking [1], though population-wide chemoprevention and active case detection are increasingly deployed in settings approaching elimination [19]. Asymptomatic and low-density infections are therefore systematically missed by routine surveillance despite contributing substantially to transmission [8, 20, 21]. Because HbAS carriers have the aforementioned reduced symptom expression while remaining parasitaemic, symptom-based testing disproportionately misses this subpopulation. This hidden reservoir is a structural consequence of biased observation. When a host subpopulation varies systematically in symptom expression while remaining infectious, under-detection is an unavoidable outcome of symptom-triggered testing. Several studies have quantified genotype-conditional infection and symptom rates [8, 11, 22], but none has directly measured the fraction of HbAS infections missed by a symptom-triggered surveillance programme in a real-world endemic setting — the operationally relevant quantity for control policy.

To explore the consequences of this interaction, we use a rule-based stochastic model implemented in the Kappa language [23]. Rule-based models represent individual entities and their interactions through explicit rewriting rules, naturally accommodating heterogeneity in host attributes — here, sickle cell status, infection state, and test results — without requiring separate compartments for every combination of states. Originally developed for molecular biology [24, 25], this formalism has been extended to epidemiological modelling [26–28] but has not, to our knowledge, been applied to malaria transmission. Its transparency — the model source closely mirrors the mathematical structure and is provided as supplementary material — supports inspection and verification of the mechanisms driving our results. While the existence of asymptomatic reservoirs and the limitations of symptom-driven surveillance are individually well established [29, 30], and Shim et al. [31] have modelled malaria with sickle cell trait to study evolutionary dynamics of the S-gene, no study has formally characterised the interaction between host genetic heterogeneity and surveillance bias or its consequences for epidemic control.

This study treats sickle cell trait not as a proven reservoir but as a biologically grounded exemplar of host heterogeneity capable of generating hidden transmission under biased surveillance. We construct a double SEI (susceptible–exposed–infectious) model for human and mosquito populations, elaborate it with sickle cell trait and a testing-treatment intervention, and conduct a systematic parameter sweep across reservoir size, test bias, and symptom specificity. The results demonstrate that the interaction between genetic reservoirs and symptom-biased testing is multiplicative: neither alone prevents epidemic control, but together they create a regime where standard surveillance strategies fail at realistic HbAS prevalence levels.

2. Materials and methods

2.1. **Rule-based modelling.** We use the Kappa language [23], a rule-based modelling formalism originally developed for molecular biology [24, 25] and subsequently applied to epidemiological modelling, where it has been used to construct a tutorial series of seven canonical models [26], a multi-scale model of immune response and virus transmission [27], and a network model of SARS-CoV-2 transmission incorporating individual-level survey data [28]. In Kappa, a system is described as a collection of agents, each possessing internal sites that can take discrete states. Dynamics are specified as rewriting rules: each rule matches a pattern of agents in particular states and transforms them, at a specified rate. This formalism is well suited to our problem because it naturally represents individual-level heterogeneity — here, the distinction between hosts with and without sickle cell trait — without requiring separate compartments for every combination of attributes. Rules are modular: the core transmission model can be elaborated with interventions (nets, testing, treatment) by adding or modifying rules without restructuring the entire model. To illustrate the practical advantage: the Kappa source for the present model is approximately 150 lines; the equivalent ODE system generated by the KaDE tool [32] is a

20-dimensional system requiring approximately 1,350 lines of Matlab code. The complete Kappa source for the model described below is provided as supplementary material.

2.2. **Host-parasite rule-based malaria model.** We develop a simple rule-based malaria transmission model that captures heterogeneity in host genetics (specifically sickle-cell trait) and the effect of interventions. The model consists of two interacting populations: humans and female mosquitoes, using a SEI model for both. Consequently, the core of the model has 6 state transitions (3 for each population: $S \to E$, $E \to I$, and $I \to S$), and the elaboration of this core model only changes the rates of these transitions based on the sickle-cell trait and the various interventions. We use fixed population sizes; specifically, we ignore seasonal or other changes in the mosquito population and define the mosquito population relative to the human population via a parameter controlling mosquitoes per person. Recovering ($I \to S$) in mosquitoes represents the mosquito dying and being replaced by a new, uninfected insect.

We introduce the models in two layers: first the core model, which introduces the six transition rules of the double SEI model, which we later elaborate with additional rules and rate changes including the sickle cell trait and the interventions.

The core model has 6 transition rules, one for each state transition.

2.2.1. *Humans.*

(1) 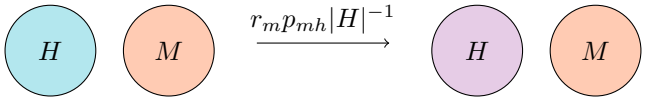

where r_m is the rate at which a single mosquito bites a human and p_{mh} is the probability that that bite results in infection. The normalisation constant, $|H|$ is the (possibly time-dependant) size of the human population. This normalisation constant arises from the computation of the biting rate for a human. The rate at which a human is bitten can be computed from the rate for a single mosquito bites, r_m, scaled by the ratio of mosquitoes to humans ($\frac{|M|}{|H|}$). The rule applies to all susceptible human and infected mosquito pairs, therefore, to get the rate for the rule, we need divide the rate calculated for a single susceptible human by the number of infected mosquitoes ($|M_I|$).

(2) 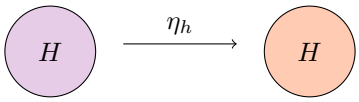

The rate of this rule, η_h is the reciprocal of the incubation period of malaria in humans. This is the standard disease progression rule.

(3) 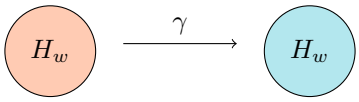

The rate of recovery in humans, η_h is the reciprocal of the average recovery period in humans. We assume *a priori* that people with sickle cell trait never recover on their own, but carry the disease until cured. We probe this assumption in the sensitivity analysis section below.

2.2.2. *Mosquitoes.*

(4) 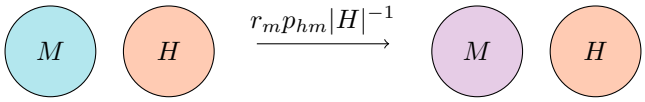

This rule is the mirror image of the rule for transmitting malaria from mosquitoes to humans and the rate is computed analogously.

(5) 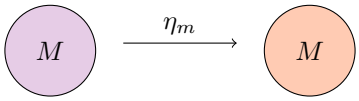

Similarly, disease progression in mosquitoes proceeds at the rate η_m

(6) 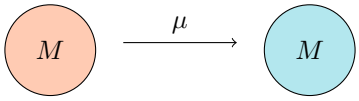

Here we have a death process for mosquitoes happening at rate μ, the reciprocal of the mosquito lifespan. In the model it is encoded as though it is recovery because we assume a constant mosquito population. A mosquito dies and is replaced with a new, susceptible mosquito.

The behaviour of the core model is investigated by changing the rates for the rules or for some sub-population (replacing the rule with multiple, mutually exclusive rules with different rates). The base rates are modified according to the three interventions the model is parametrised by: mosquito population control, nets, and testing-treatment. Mosquito population control directly reduces the number of female mosquitoes per person. Nets are modelled by reducing the chance that a female mosquito bites a human. We differentiate between people with sickle-cell trait and those who do not have it. Using two parameters based on host-genetics allows us to explore the assumption that people with sickle-cell trait, who never experienced the symptoms of malaria, are less likely to use nets.

Modelling testing and treatment is more involved, as we allow a combination of different testing strategies. For simplicity we assume that tests are 100% accurate. We assume that anyone who gets a positive test gets treatment, which reduces the duration of malaria infection. We model this by changing the rate for infection rule for humans with positive test to the reciprocal of the average duration of malaria with treatment.

The model allows the combination (and thus inspecting the effectiveness) of two testing strategies: uniform selection and symptomatic selection. The first tests people uniformly independent of any conditions, while the latter only selects people with symptoms (but otherwise uniformly). This second testing strategy is more complex to implement, as this is the point in the model where we take into consideration that people with sickle-cell trait do not show symptoms. We also introduce a symptom specificity parameter, σ_o, representing the probability that an uninfected individual does *not* present malaria-like symptoms. When specificity is 1, only truly infected non-sickle-cell individuals are symptomatic; when specificity is less than 1, a fraction of uninfected individuals present with symptoms indistinguishable from malaria (due to other febrile illnesses or non-specific presentations). This parameter controls the false positive rate of symptom-driven surveillance without requiring an explicit co-infection model: the effective test bias is diluted when some symptomatic individuals turn out not to have malaria.

2.3. **Experimental design.** To investigate the interaction between hidden genetic reservoirs and testing strategy, we performed a systematic parameter sweep across three dimensions. The sickle cell prevalence S (reservoir size) was varied from 0% to 40% in increments of 5%, spanning the range from a control with no genetic reservoir to approximately the maximum HbAS prevalence observed in highly endemic regions of Sub-Saharan Africa. The test bias parameter was varied from 0 (purely random testing, where individuals are selected for testing uniformly regardless of symptoms) to 1 (purely symptomatic testing, where only individuals presenting symptoms are tested) in increments of 0.05. The test rate was varied across the range shown in the figures, controlling the overall volume of testing effort.

For each combination of reservoir size and test bias, the model was simulated for 1000 days or until the epidemic was extinguished (i.e. no infected humans or mosquitoes remained), whichever occurred first. Two outcome measures were recorded: the average fraction of humans infected at the end of the simulation (InfH_last_avg), capturing equilibrium prevalence, and the cumulative number of tests administered (Tests_last_avg), capturing the cost of each testing strategy.

The full sweep was conducted at two levels of symptom specificity: $\sigma_o = 1.00$ (perfect specificity, where only truly infected non-sickle-cell individuals present symptoms) and $\sigma_o = 0.99$ (1% false positive rate, where 1% of uninfected individuals also present malaria-like symptoms). This yields four series of

Symbol	Description	Value	Reference
r_m	Mosquito biting rate	0.5 day^{-1}	[35]
p_{mh}	Probability of transmission, mosquito to human	0.09	[37]
p_{hm}	Probability of transmission, human to mosquito	0.09	[38]
η_h^{-1}	Human incubation period	10.5 days	[39]
η_m^{-1}	Mosquito incubation period	10 days	[39]
γ^{-1}	Human recovery rate (untreated, non-HbAS)	200 days	—
μ^{-1}	Mosquito lifespan	42 days	[40]
τ^{-1}	Recovery rate after positive test	1.6 days	[41]
m	Female mosquitoes per human	0.6	[42]
S	Sickle cell prevalence	0–0.40 (swept)	[34]
d	Overall testing rate	0–0.0095 tests per person per day (swept)	—
β	Fraction of testing allocated to symptomatic	0–1 (swept)	—
σ_o	Symptom specificity	1.00 or 0.99	—
N	Human population	10,000	—

TABLE 1. Model parameters. Fixed parameters are drawn from the literature where available. Swept parameters define the experimental design described in Section 3.1.

heatmaps: prevalence at perfect specificity, prevalence at 99% specificity, test count at perfect specificity, and test count at 99% specificity — each comprising 9 heatmaps (one per reservoir size), for a total of 36 figures.

All other model parameters were held constant at the values described in Section 3.1. The human population was set to $N = 10{,}000$ for computational tractability, with initial infection rates of 50% in both humans and mosquitoes. Simulations were performed using KaSim [33].

3. RESULTS

3.1. **Parameter values.** We obtain most of our malaria model parameters directly from the data in the literature. The prevalence of sickle cell trait (heterozygous carrier status HbAS) in Nigeria was assumed to be 25% (0.25) as in [34]. The mosquito biting rate is assumed to be 0.5 per day as in [35], who used that value based on [36]. The probability of transmission from mosquito to human is about 9%, ranging from 4.5% to 16% [37], giving $p_{mh} = 0.09$. The probability of transmission from human to mosquito depends on the prevalence of malaria in the population and fluctuates between 5% and 14% [38]; we adopt $p_{hm} = 0.09$. The incubation period for malaria (infection with *P. falciparum*) is 7 to 14 days [39]; we assume an average of $\eta_h^{-1} = 10.5$ days. The mosquito lifespan is assumed to be $\mu^{-1} = 42$ days as in [40]. The incubation period in mosquitoes is 9 days or longer [39]; we assume $\eta_m^{-1} = 10$ days. The recovery rate for untreated, non-sickle-cell humans is set to $\gamma^{-1} = 200$ days, corresponding to an average infection duration of approximately 200 days in the absence of treatment. The treatment rate for individuals who test positive is set to $\tau^{-1} = 1.6$ days, following [41]. The ratio of female *Anopheles* mosquitoes to humans is set to $m = 0.6$. Net effectiveness is modelled as a 50% reduction in biting rate.

3.2. **Reservoir size and test bias jointly determine epidemic control.** Figure 1 shows the average fraction of humans infected at equilibrium (or at 1000 days) as a function of test bias (horizontal axis) and test rate (vertical axis), for three representative reservoir sizes: $S = 0$ (no sickle cell carriers, the control), $S = 0.20$ (approximate HbAS prevalence in Nigeria), and $S = 0.40$ (the upper bound of the range examined).

In the control case ($S = 0$), a sharp threshold separates the suppressed and endemic regimes. Above a test rate of approximately 0.002, the epidemic is controlled regardless of test bias. Below this threshold, test bias matters: purely symptomatic testing ($\beta = 1$) leaves residual infection at low test rates, while random

testing ($\beta = 0$) remains effective at modestly lower rates. The overall picture is that in a population without a genetic reservoir, testing at even moderate rates is sufficient to suppress the epidemic, and the choice between random and symptomatic testing is of secondary importance.

As reservoir size increases, this picture changes qualitatively. At $S = 0.20$, the transition between suppressed and endemic regimes is no longer a flat threshold but a diagonal: the test rate required for suppression increases with test bias. Under purely symptomatic testing, the epidemic persists at test rates that would be more than sufficient under random testing. The effect is approximately a doubling of the required testing effort — random testing at a rate of 0.002 achieves comparable suppression to symptomatic testing at a rate of 0.004 or higher. This is the operationally relevant case: it corresponds to real-world HbAS prevalence in the regions most affected by *P. falciparum* malaria.

At $S = 0.40$, the endemic region dominates the lower half of the parameter space. The transition boundary has shifted upward across the entire range of test bias, and the boundary itself has broadened from a sharp threshold to a diffuse gradient. Under high test bias, substantial infection persists even at the highest test rates examined. The bottom-right corner of the heatmap — high bias, low test rate — shows equilibrium infection levels of 0.5 or higher. At this reservoir size, symptom-biased testing is essentially unable to control the epidemic regardless of intensity.

The full sequence of nine heatmaps (0% to 40% in 5% increments, presented in the Supplementary Material) reveals a smooth and monotonic progression between these extremes. Several features are notable. First, the effect is already visible at 10–15% reservoir size, well within the range of HbAS prevalence observed across Sub-Saharan Africa. Second, the interaction between reservoir size and test bias is multiplicative rather than additive: neither parameter alone produces epidemic failure, but their combination creates a regime where testing cannot compensate for bias. Third, even at zero test bias, the transition threshold shifts upward as reservoir size increases, reflecting the baseline cost that a hidden reservoir imposes on any testing strategy — random testing still works, but requires more effort.

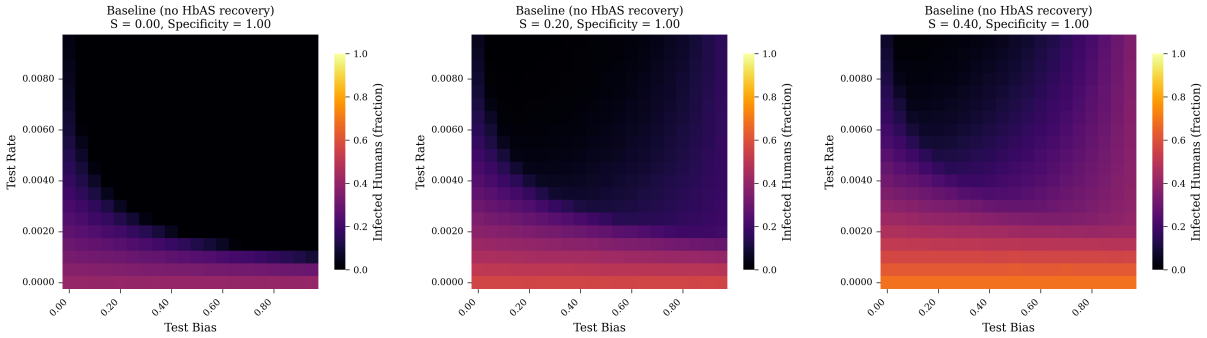

(A) $S = 0$ (control) (B) $S = 0.20$ (C) $S = 0.40$

FIGURE 1. Equilibrium infection prevalence as a function of test bias and test rate, for three reservoir sizes. (a) No sickle cell carriers (control). (b) Approximate Nigerian HbAS prevalence. (c) Upper bound of the range examined. Colour scale indicates the average fraction of humans infected at the end of the 1000-day simulation. Dark regions indicate epidemic suppression; bright regions indicate persistent endemic infection. Symptom specificity is 1.00 (perfect).

3.3. **Symptom specificity dilutes test bias.** Figure 2 compares the prevalence heatmaps at $S = 0.20$ under perfect symptom specificity (obs_specificity = 1.00) and 99% specificity (obs_specificity = 0.99). When 1% of uninfected individuals present with malaria-like symptoms, the outcomes under high test bias improve slightly: the transition boundary sits fractionally lower, and the bottom-right corner is marginally less intense.

This effect is counterintuitive. The naive expectation is that false positives in symptom recognition should waste testing capacity and worsen outcomes. The actual mechanism is different: when some non-infected individuals present symptoms, selecting symptomatic individuals for testing becomes slightly more like random testing. The false positives dilute the effective test bias, pulling it back towards zero.

Since the fundamental problem is the bias itself — not the volume or accuracy of testing — this dilution is beneficial.

The effect is most visible in the mid-range reservoir sizes (15–25%), where the system is in the transition zone between controlled and uncontrolled epidemic regimes. At $S = 0$, there is little bias effect to dilute, so the two specificity levels produce nearly identical results. At $S = 0.40$, the reservoir dominates so strongly that a 1% dilution barely registers. At zero test bias, the two series are identical regardless of reservoir size, as expected: random testing is unaffected by symptom specificity.

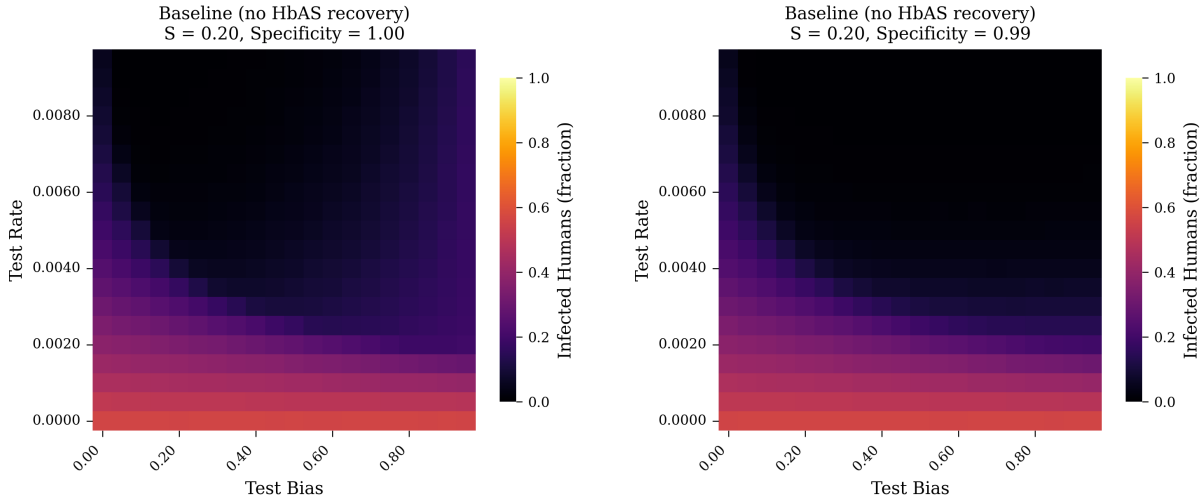

(A) Specificity = 1.00 (B) Specificity = 0.99

FIGURE 2. Effect of symptom specificity on epidemic control at $S = 0.20$. (a) Perfect specificity (obs_specificity = 1.00). (b) 99% specificity (obs_specificity = 0.99). The slight improvement in the high-bias region under imperfect specificity reflects dilution of effective test bias by false positive symptoms.

3.4. **Symptom-biased testing is a false economy.** The cumulative number of tests administered over each simulation provides a complementary perspective on the cost of different testing strategies. Figure 3 shows the test count heatmaps for $S = 0$ and $S = 0.40$ at perfect specificity.

The gradient in test counts runs from top-left (high test rate, low bias: many tests administered, up to ∼100,000) to bottom-right (low test rate, high bias: few tests administered). This is the mirror image of the prevalence heatmaps: the region of parameter space where fewest tests are used is precisely the region where the epidemic is least well controlled.

Crucially, the test count heatmaps are nearly invariant to reservoir size. The number of tests consumed is determined by the testing policy (rate and bias), not by the amount of disease present. This means that symptom-biased testing costs the same whether or not it works. At $S = 0$, biased testing is cheap and reasonably effective — a defensible policy choice. At $S = 0.40$, biased testing is still cheap but completely fails to control the epidemic. The cost of random testing is the same in both cases, but its value increases dramatically as the reservoir grows.

Under 99% symptom specificity, the test count gradient becomes slightly smoother: false positive symptoms generate a small number of additional tests under high bias, representing the cost side of the dilution effect described above. At $S = 0.40$ with 99% specificity, bias barely changes the test count — the false positives have substantially washed out the distinction between symptomatic and random testing in terms of both cost and effectiveness.

The policy implication is direct. Symptom-biased testing reduces test volume but fails to control the epidemic in populations with significant HbAS prevalence. It is a false economy: the savings in test kits are purchased at the cost of sustained endemic transmission.

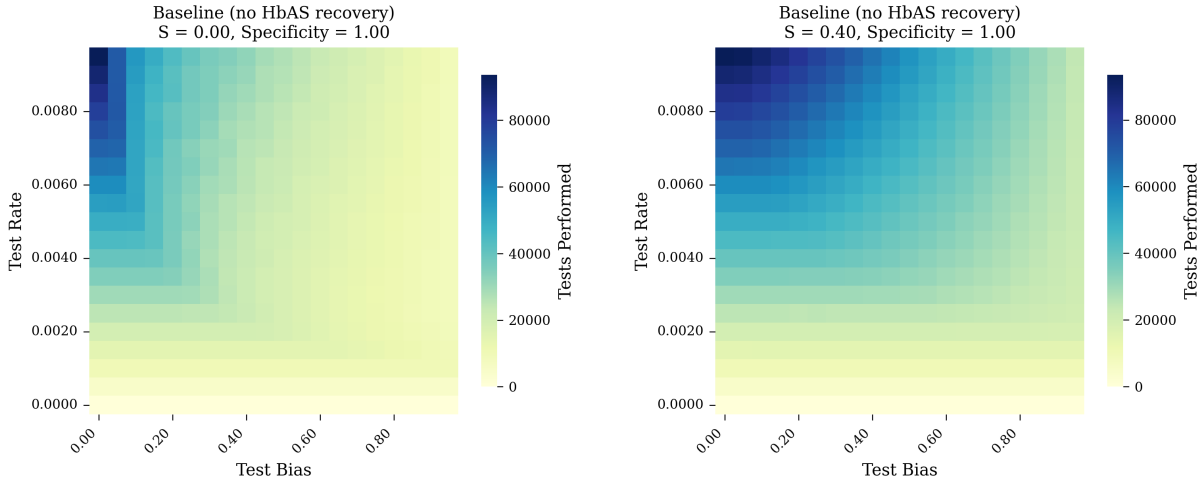

(A) $S = 0$ (control) (B) $S = 0.40$

FIGURE 3. Cumulative tests administered as a function of test bias and test rate. (a) No sickle cell carriers (control). (b) $S = 0.40$. Colour scale indicates total tests over the simulation. The gradient from top-left (many tests) to bottom-right (few tests) is the mirror image of the prevalence heatmaps: the cheapest testing strategies are the least effective. Test counts are nearly invariant to reservoir size — the cost of a testing policy is determined by the policy, not by its effectiveness.

3.5. **Sensitivity to HbAS recovery rate.** The baseline model assumes that HbAS carriers never clear infection spontaneously — they must be identified and treated. This is a strong assumption, and it is not universally accepted that sickle cell carriers recover more slowly than non-carriers. To assess whether the main results depend on this assumption, we repeated the full parameter sweep under two alternative recovery scenarios: one in which HbAS carriers recover at the same rate as non-carriers ($\gamma_{\text{HbAS}} = 1/200$ day^{-1}, equal recovery), and one in which they recover at half the rate ($\gamma_{\text{HbAS}} = 1/400$ day^{-1}, slow recovery).

Figure 4 compares the three recovery assumptions at two reservoir sizes: $S = 0.20$ (top row) and $S = 0.40$ (bottom row), both at perfect specificity.

At $S = 0.20$, the qualitative pattern is preserved across all three scenarios: the diagonal transition boundary between suppressed and endemic regimes remains, and symptom-biased testing continues to require substantially higher test rates than random testing. The boundary shifts downward as the HbAS recovery rate increases — equal recovery produces a slightly more favourable landscape than no recovery — but the shift is modest.

The differences become substantially more visible at $S = 0.40$, where the recovery assumption has a clear quantitative effect. Under equal recovery, a large suppressed region persists in the upper-left of the parameter space — testing can still control the epidemic at sufficiently high rates and low bias. Under no recovery (the baseline), the endemic region dominates and the suppressed region retreats to only the highest test rates. Slow recovery is intermediate. Nevertheless, the diagonal structure of the transition boundary is present in all three cases: even under equal recovery at $S = 0.40$, symptom-biased testing fails at rates where random testing succeeds.

This is the central point of the sensitivity analysis: the hidden reservoir effect is driven primarily by observation bias (HbAS carriers are asymptomatic and therefore invisible to symptom-driven testing), not by differential recovery. Slower recovery amplifies the effect by keeping HbAS carriers infectious for longer, but even without any recovery differential, the structural mismatch between where the disease is and where the testing looks is sufficient to sustain endemic transmission under biased surveillance.

The results are also insensitive to initial conditions. We repeated the baseline sweep at $S = 0.20$ with initial infection rates of 10% and 90% (compared with the default 50%) in both humans and mosquitoes.

The equilibrium heatmaps are virtually indistinguishable across all three initial conditions. This is expected: the outcome measure is the equilibrium (or 1000-day) prevalence, so the system has time to converge regardless of starting point. The insensitivity to initial conditions confirms that the results reflect stable equilibrium behaviour rather than transient dynamics dependent on how the epidemic begins.

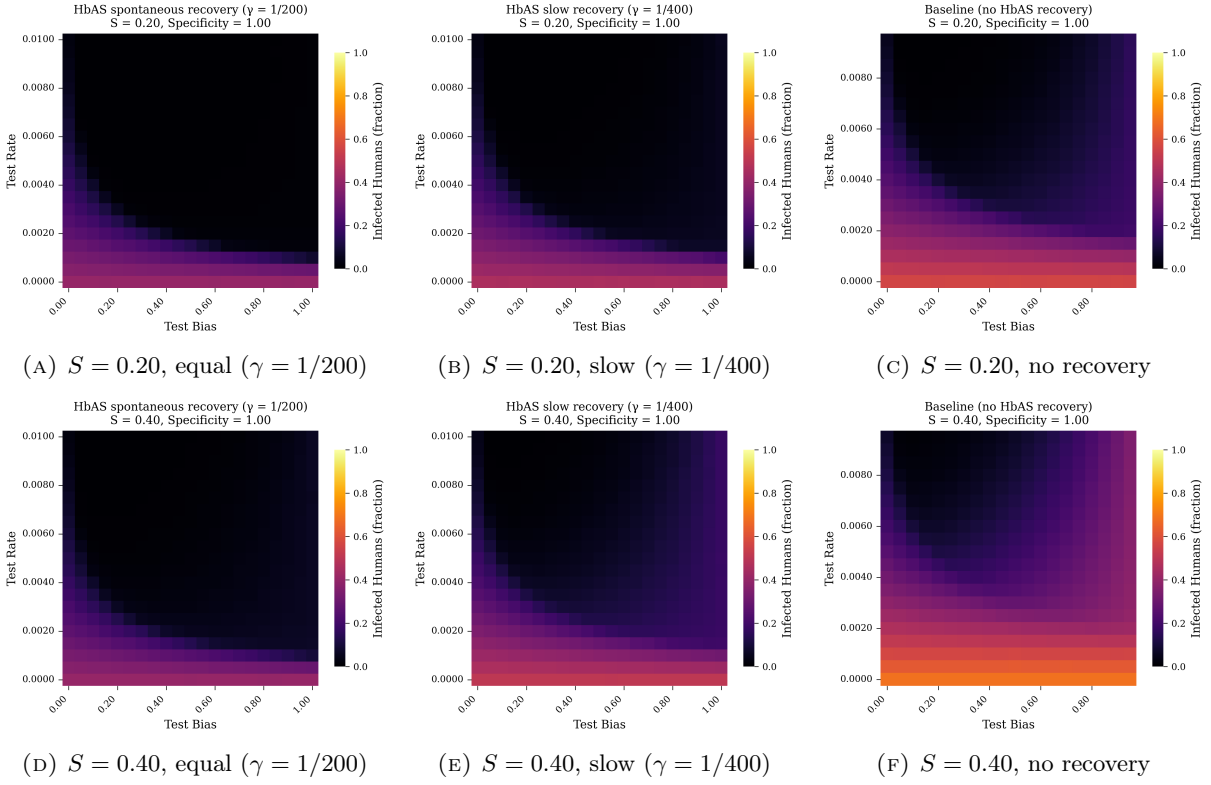

(A) $S = 0.20$, equal ($\gamma = 1/200$) (B) $S = 0.20$, slow ($\gamma = 1/400$) (C) $S = 0.20$, no recovery

(D) $S = 0.40$, equal ($\gamma = 1/200$) (E) $S = 0.40$, slow ($\gamma = 1/400$) (F) $S = 0.40$, no recovery

FIGURE 4. Sensitivity of epidemic control to the HbAS recovery rate. Top row: $S = 0.20$. Bottom row: $S = 0.40$. Columns from left to right: HbAS carriers recover at the same rate as non-carriers ($\gamma = 1/200$), at half the rate ($\gamma = 1/400$), or never (baseline). The diagonal transition boundary persists across all scenarios; the recovery assumption affects the quantitative position of the boundary — particularly at high reservoir sizes — but not the qualitative pattern.

4. DISCUSSION

The central finding of this study is that the interaction between host genetic heterogeneity and symptom-biased surveillance creates a regime in which standard testing strategies systematically fail to control malaria transmission. This result is structural: it emerges from the logic of biased observation under heterogeneous symptom expression, rather than from any single parameter value or empirical assumption about sickle cell trait. The progressive degradation of epidemic control across the reservoir size sweep — from sharp, easily achievable thresholds at 0% HbAS prevalence to diffuse, unreachable boundaries at 40% — demonstrates that the phenomenon is robust and monotonic, not a fragile artefact of model calibration.

The rule-based modelling framework [23] is well suited to this problem precisely because it represents individual-level heterogeneity natively. Unlike compartmental models that aggregate over population subgroups, the Kappa formalism allows each agent to carry its own combination of attributes — sickle cell status, infection state, test result — and the interaction rules operate on these attributes directly. This transparency is important: the model's behaviour can be inspected rule by rule, and the mechanisms driving the results are legible rather than buried in differential equations. The complete model source is provided as supplementary material to support this inspection.

4.1. **Policy implications.** The practical implication is direct. In populations where HbAS prevalence is 20–25% — which includes Nigeria and much of the West African malaria belt — symptom-driven testing strategies require roughly double the testing effort to achieve the same level of epidemic suppression as random testing. At higher reservoir sizes, no feasible level of symptom-biased testing achieves suppression at all. Since symptom-driven testing is the default mode of malaria surveillance in most endemic settings, this represents a structural blind spot in current control strategies.

The test count analysis reinforces this point. Symptom-biased testing uses fewer test kits than random testing — it is cheaper per test administered. But it fails to control the epidemic, meaning the investment in testing yields poor returns. Random testing costs more in absolute terms but succeeds in suppressing transmission. The implication for resource allocation is that the efficiency of a testing strategy should be evaluated not by cost per test but by cost per unit of epidemic suppression. Viewed this way, symptom-biased testing in high-HbAS populations is a false economy.

The parameter sweep allows us to quantify the benefit of incorporating random testing. At $S = 0.20$ and perfect specificity, purely symptomatic testing (test_bias = 0.95) never suppresses the epidemic within the simulated range: even at the highest test rate examined (0.0095 per person per day), equilibrium prevalence remains above 15%. A strategy allocating 80% of tests symptomatically and 20% randomly (test_bias = 0.80) improves matters but still fails to suppress, reaching a floor of approximately 7% prevalence. Shifting to a 50/50 allocation (test_bias = 0.50) crosses the suppression threshold at a test rate of 0.0055, consuming approximately 38,000 cumulative tests over the 1000-day simulation. A strategy allocating 80% of tests randomly and 20% symptomatically (test_bias = 0.20) achieves suppression at the lower rate of 0.0045, with approximately 44,000 tests. Notably, purely random testing (test_bias = 0) is marginally less efficient than these mixed strategies: it only just fails to suppress at the highest rate examined (InfH = 0.051 at rate = 0.0095, approximately 94,000 tests), because it foregoes the advantage of cheaply catching the symptomatic infections that are easy to find. The optimal regime is thus a hybrid: predominantly random testing supplemented by a symptomatic component, rather than either pure strategy alone (Table 2).

Strategy	**Threshold rate**	**Tests at threshold**	**Equilibrium InfH**
95% symptomatic	never	21,456†	0.152
80% symptomatic	never	32,024†	0.073
50% symptomatic	0.0055	38,132	0.046
80% random	0.0045	43,973	0.045
90% random	0.0060	59,021	0.024
100% random	never	93,519†	0.051

†At highest simulated rate (0.0095). "Never" = not suppressed within range.

TABLE 2. Suppression thresholds by testing strategy at $S = 0.20$, obs_specificity = 1.00. Threshold is the minimum test rate at which equilibrium prevalence falls below 5%. Tests are cumulative over the 1000-day simulation for a population of $N = 10{,}000$.

These results suggest that existing surveillance programmes in high-HbAS regions could substantially improve epidemic control by reallocating a fraction of testing capacity from symptom-triggered to random screening. A hybrid strategy allocating roughly half of tests randomly achieves suppression at moderate test rates, while purely symptomatic testing fails regardless of intensity.

4.2. **The role of symptom specificity.** The comparison between perfect and imperfect symptom specificity reveals a counterintuitive effect: a 1% false positive rate in symptom recognition slightly *improves* epidemic control under biased testing. The mechanism is that false positives dilute the effective test bias — when some non-infected symptomatic individuals are tested, the testing strategy becomes slightly more random than intended. This dilution is beneficial precisely because the fundamental problem is the bias itself, not the accuracy of symptom recognition.

This finding has implications for how we think about diagnostic specificity in surveillance contexts. In clinical medicine, false positives are unambiguously costly: they lead to unnecessary treatment, patient anxiety, and wasted resources. In population-level surveillance aimed at epidemic control, the calculus is

different. A test that occasionally flags non-infected individuals may inadvertently broaden the surveillance net, catching infections that would otherwise be missed by a purely symptom-targeted strategy. This does not argue for deliberately reducing diagnostic specificity, but it does suggest that the epidemiological cost of imperfect specificity in symptom-driven testing may be lower than expected — and in some regimes, the net effect may be positive.

4.3. **Limitations.** Several limitations of the present study should be acknowledged. First, the model assumes that tests are perfectly accurate: any infected individual who is tested receives a positive result and subsequent treatment. It is important to distinguish this assumption from the observation bias that drives our main result. The hidden reservoir effect arises because HbAS carriers are asymptomatic and therefore *never selected for testing* under symptom-biased surveillance — not because tests fail to detect their infections. Perfect test accuracy is a best-case assumption for symptom-biased testing: it means that the failure of symptomatic surveillance documented here is entirely due to who gets tested, not to whether the test works. In practice, rapid diagnostic tests and microscopy have well-documented sensitivity limitations, particularly at the low parasite densities characteristic of asymptomatic and submicroscopic infections [43]. Since HbAS carriers tend to harbour lower parasite densities, imperfect test sensitivity would compound the observation bias by missing some HbAS infections even among carriers who are tested. The results presented here are therefore a lower bound on the true magnitude of the hidden reservoir problem: relaxing the perfect-test assumption would worsen outcomes for all strategies, but disproportionately for those that rely on finding asymptomatic carriers through random screening.

Second, the model does not capture treatment-seeking behaviour. Not all symptomatic individuals present to health facilities, and the probability of seeking care may itself depend on factors correlated with sickle cell status, distance to clinics, cost of testing, and cultural attitudes towards malaria. The introduction of user fees for rapid diagnostic tests, replacing previously free provision, has been shown to reduce testing uptake in some settings. Our model implicitly assumes that all symptomatic individuals are available for testing, which overestimates the reach of symptom-biased surveillance.

Third, the model uses a fixed population size and does not account for seasonal variation in mosquito populations, migration, or demographic change. The simulations use a population of $N = 10{,}000$ for computational tractability, which is sufficient to capture the qualitative dynamics of interest but may not reproduce the stochastic effects that become important in very small or very large populations. The sensitivity analysis in Section 3.5 confirms that the results are insensitive to initial conditions (10%, 50%, and 90% initial infection all converge to the same equilibrium), consistent with the fact that the outcome measure reflects long-run behaviour.

Fourth, the baseline recovery assumption for HbAS carriers — that they never clear infection spontaneously — is an extreme case. The sensitivity analysis in Section 3.5 demonstrates that the qualitative results are robust to this assumption: even when HbAS carriers recover at the same rate as non-carriers, the hidden reservoir effect persists. Nevertheless, the true recovery dynamics of HbAS carriers remain poorly characterised, and field data on spontaneous parasite clearance rates conditional on sickle cell genotype would be valuable for calibrating future models.

Finally, although we frame the analysis in terms of sickle cell trait and *P. falciparum* malaria, the underlying mechanism is general. Any infectious disease in which a genetically or otherwise defined subpopulation experiences attenuated symptoms while remaining infectious will exhibit analogous hidden reservoir dynamics under symptom-biased surveillance. The specific parameters will differ, but the structural interaction between host heterogeneity and observation bias is not disease-specific.

4.4. **Relationship to prior work.** The rule-based modelling approach used here builds on the Kappa framework developed for molecular biology [24, 25] and extended to epidemiological applications. To our knowledge, this is the first application of rule-based modelling to malaria transmission dynamics, and specifically to the question of how host genetic heterogeneity interacts with surveillance bias. The compositional nature of the formalism — where the core transmission model can be elaborated with interventions by adding rules rather than restructuring the model — proved particularly valuable for exploring the interaction between biological and observational mechanisms.

The closest prior work is that of Shim et al. [31], who constructed a compartmental model of malaria incorporating sickle cell trait and distinguishing symptomatic from asymptomatic infections. Their biological premise — that HbAS protects against symptomatic disease but not infection — is the same as ours. However, their focus was on the evolutionary dynamics of the sickle cell gene: their key finding was that controlling symptomatic malaria may reduce S-gene frequency, eventually increasing overall prevalence. They did not examine testing strategy as a variable, did not sweep over surveillance parameters, and did not identify the multiplicative interaction between reservoir size and observation bias that we characterise here.

More broadly, Chisholm et al. [29] reviewed the implications of asymptomatic carriers for infectious disease transmission models and showed that ignoring carriers leads to overestimated transmission rates and incorrectly assessed interventions. Their analysis established the general principle that carrier states matter for model validity, but addressed model misspecification rather than surveillance strategy design. Park and Bolker [30] demonstrated that incorrect assumptions about observation processes can bias estimates of the basic reproduction number, formalising the statistical consequences of imperfect surveillance without examining how surveillance design interacts with host heterogeneity. Agent-based malaria models such as EMOD [44] have characterised the infectious reservoir by age and detectability, showing that interventions shift the reservoir towards sub-microscopic infections, but their focus is on diagnostic sensitivity rather than on the structural interaction between host genetics and testing strategy.

Our contribution bridges these two literatures. The empirical literature documents that asymptomatic reservoirs exist and sustain transmission; the theoretical literature establishes that observation bias matters for parameter estimation and model validity. What has been missing is a formal characterisation of how these two factors interact to determine epidemic control outcomes — specifically, the shape of the failure regime, the multiplicative rather than additive nature of the interaction, and the operationally relevant parameter values at which surveillance breaks down. That is what the present study provides.

5. Conclusion

Symptom-driven malaria surveillance systematically under-detects infections in populations carrying sickle cell trait. Using a rule-based stochastic model, we have shown that this under-detection is not merely a measurement error but a structural feature of biased testing under host heterogeneity — one that can sustain endemic transmission even when testing effort is high. At realistic HbAS prevalence levels, the interaction between hidden genetic reservoirs and symptom-biased testing creates a regime where standard surveillance strategies fail. Incorporating random testing into malaria surveillance programmes in HbAS-prevalent regions could substantially improve epidemic control at modest additional cost. More broadly, these results highlight the need to account for host genetic heterogeneity when designing and evaluating infectious disease surveillance strategies.

Software and Data Availability. The code and summary data underlying this publication are available at `https://codeberg.org/rbem/malaria`. Full datasets of all simulation output is available upon request.

Author Contributions. SB formulated the model, and conducted the numerical experiments. AAO originated the idea of investigating the sickle cell reservoir and provided parameter values. DS contributed parameter values and contextual framing. WW contributed funding and strategic direction. All authors drafted the text.

Acknowledgements. This work was supported by the UK MRC-NIHR Better Methods Better Research Grant X/011658/1.

References

[1] World Health Organization. World malaria report 2025. Technical report, World Health Organization, Geneva, December 2025. URL https://www.who.int/teams/global-malaria-programme/reports/world-malaria-report-2025. Accessed: 2026-03-09.

[2] UNICEF Nigeria. Nigeria receives malaria vaccines ahead of roll out, October 2024. URL https://www.unicef.org/nigeria/press-releases/nigeria-receives-malaria-vaccines-ahead-roll-out. Accessed: 2026-03-09.

[3] Gavi, the Vaccine Alliance. Malaria vaccine support, January 2026. URL https://www.gavi.org/types-support/vaccine-support/malaria. Accessed: 2026-03-09.

[4] Carolina M Andrade, Hannah Fleckenstein, Richard Thomson-Luque, Safiatou Doumbo, Nathalia F Lima, Carrie Anderson, Julia Hibbert, Christine S Hopp, Tuan M Tran, Shanping Li, et al. Increased circulation time of Plasmodium falciparum underlies persistent asymptomatic infection in the dry season. *Nature Medicine*, 26:1929–1940, 2020. doi:10.1038/s41591-020-1084-0.

[5] Teun Bousema, Lucy Okell, Ingrid Felger, and Chris Drakeley. Asymptomatic malaria infections: detectability, transmissibility and public health relevance. *Nature Reviews Microbiology*, 12:833–840, 2014. doi:10.1038/nrmicro3364.

[6] Kimberly M Fornace, Nor Azizah Nuin, Tock H Chua, Timothy William, and Chris Drakeley. Epidemiology of *Plasmodium knowlesi* malaria, 2021.

[7] Teun Bousema and Chris Drakeley. Epidemiology and infectivity of Plasmodium falciparum and Plasmodium vivax gametocytes in relation to malaria control and elimination. *Clinical Microbiology Reviews*, 24:377–410, 2011. doi:10.1128/CMR.00051-10.

[8] Chiara Andolina, John C Rek, Jessica Briggs, Joseph Okoth, Alex Musiime, Jordache Ramjith, Noam Teyssier, Melissa Conrad, Joaniter I Nankabirwa, Kjerstin Lanke, Isabel Rodriguez-Barraquer, Lisette Meerstein-Kessel, Emmanuel Arinaitwe, Peter Olwoch, Philip J Rosenthal, Moses R Kamya, Grant Dorsey, Bryan Greenhouse, Chris Drakeley, Sarah G Staedke, and Teun Bousema. Sources of persistent malaria transmission in a setting with effective malaria control in eastern Uganda: a longitudinal, observational cohort study. *Lancet Infectious Diseases*, 21:1568–1578, 2021. doi:10.1016/S1473-3099(21)00072-4.

[9] Helena Lamptey, Michael F Ofori, Bright Adu, et al. Impact of haemoglobinopathies on asymptomatic plasmodium falciparum infection and naturally acquired immunity among children in northern ghana. *Frontiers in Hematology*, 2, 2023. doi:10.3389/frhem.2023.1150134.

[10] Mahamadou Aidoo, Dianne J Terlouw, Margarette S Kolczak, Peter D McElroy, Feiko O ter Kuile, Simon Kariuki, Bernard L Nahlen, Altaf A Lal, and Venkatachalam Udhayakumar. Protective effects of the sickle cell gene against malaria morbidity and mortality. *Lancet*, 359(9314), April 2002. doi:10.1016/S0140-6736(02)08273-9.

[11] Thomas N Williams, Tabitha W Mwangi, Sammy Wambua, Neal D Alexander, Moses Kortok, Robert W Snow, and Kevin Marsh. Sickle cell trait and the risk of plasmodium falciparum malaria and other childhood diseases. *J Infect Dis*, 192(1), 2005. doi:10.1086/430744.

[12] Liang Gong, Catherine Maiteki-Sebuguzi, Philip J Rosenthal, Alan E Hubbard, Chris J Drakeley, Grant Dorsey, and Bryan Greenhouse. Evidence for both innate and acquired mechanisms of protection from plasmodium falciparum in children with sickle cell trait. *Blood*, 119(16), 2012. doi:10.1182/blood-2011-08-371062.

[13] Helena Lamptey, Michael F Ofori, Bright Adu, Kwadwo A Kusi, Emmanuel K Dickson, Eric Kyei-Baafour, and Linda E Amoah. Association between alpha-thalassaemia trait, plasmodium falciparum asexual parasites and gametocyte carriage in a malaria endemic area in southern ghana. *BMC Research Notes*, 12, 2019. doi:10.1186/s13104-019-4181-8.

[14] Youssouf Djimdé Lawaly, Adama Sacko, Ibrahim Ag Bazet, et al. Heritability of the human infectious reservoir of malaria parasites. *PLOS ONE*, 5(6), 2010. doi:10.1371/journal.pone.0011358.

[15] Bronner P Gonçalves, Issaka Sagara, Mamadou Coulibaly, Yimin Wu, Mahamadoun H Assadou, Agnes Guindo, Ruth D Ellis, Mahamadou Diakite, Erin Gabriel, D Rebecca Prevots, Ogobara K Doumbo, and Patrick E Duffy. Hemoglobin variants shape the distribution of malaria parasites in human populations and their transmission potential. *Scientific Reports*, 7:14267, 2017. doi:10.1038/s41598-017-14627-y.

[16] Christelle Maffo Ngou, Sandrine Eveline Mandeng, Christophe Noel Noumedem Anangmo, et al. Influence of the sickle cell trait on plasmodium falciparum infectivity from naturally infected gametocyte carriers. *BMC Infectious Diseases*, 23(317), 2023. doi:10.1186/s12879-023-08134-x.

[17] World Health Organization. Guidelines for the treatment of malaria, 3rd edition, 2015. URL https://www.who.int/publications/i/item/9789241549127. Accessed: 2026-03-09.

[18] Michelle E Roh, Kelly Harvard, et al. Testing and treatment for malaria elimination: a systematic review. *Malaria Journal*, 22:254, 2023. doi:10.1186/s12936-023-04670-8.

[19] World Health Organization. Mass drug administration for falciparum malaria: A practical field manual. Technical report, World Health Organization, Geneva, 2019. URL https://www.who.int/publications/i/item/9789241515030. Accessed: 2026-03-09.

[20] Hannah C Slater, Amanda Ross, Ingrid Felger, Natalie E Hofmann, Leanne Robinson, Jackie Cook, Bronner P Gonçalves, Ottar N Bjørnstad, Felistas Mwingira, Blaise Genton, Marcel Tanner, Helena Sick, et al. The temporal dynamics and infectiousness of subpatent Plasmodium falciparum infections in relation to disease burden. *Nature Communications*, 10:1433, 2019. doi:10.1038/s41467-019-09441-1.

[21] Jean Langhorne et al. Effect of weekly fever-screening and treatment and monthly RDT testing and treatment on the infectious reservoir of malaria parasites in Burkina Faso: a cluster-randomised trial. *Lancet Microbe*, 2024. doi:10.1016/S2666-5247(24)00114-9.

[22] Catherine Kamya et al. Age-related changes in malaria clinical phenotypes during infancy are modified by sickle cell trait. *Journal of Infectious Diseases*, 224:1513–1522, 2021. doi:10.1093/infdis/jiab275.

[23] Vincent Danos and Cosimo Laneve. Formal molecular biology. *Theoretical Computer Science*, 325 (1):69–110, September 2004. ISSN 0304-3975. doi:10.1016/j.tcs.2004.03.065.

[24] Vincent Danos, Jérôme Feret, Walter Fontana, Russell Harmer, and Jean Krivine. Rule-based modelling of cellular signalling. In Luís Caires and Vasco T. Vasconcelos, editors, *CONCUR 2007 – Concurrency Theory*, pages 17–41, Berlin, Heidelberg, 2007. Springer Berlin Heidelberg. ISBN 978-3-540-74407-8.

[25] Agnes Köhler, Jean Krivine, and Jakob Vidmar. A rule-based model of base excision repair. In Pedro Mendes, Joseph O. Dada, and Kieran Smallbone, editors, *Computational Methods in Systems Biology*, pages 173–195, Cham, 2014. Springer International Publishing. ISBN 978-3-319-12982-2.

[26] W. Waites, M. Cavaliere, D. Manheim, J. Panovska-Griffiths, and V. Danos. Rule-based epidemic models. *Journal of Theoretical Biology*, 526:110851, July 2021. ISSN 0022-5193. doi:10.1016/j.jtbi.2021.110851.

[27] W. Waites, M. Cavaliere, V. Danos, R. Datta, R. M. Eggo, T. B. Hallett, D. Manheim, J. Panovska-Griffiths, T. W. Russell, and V. I. Zarnitsyna. Compositional modelling of immune response and virus transmission dynamics. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 380(2233):20210307, October 2022. ISSN 1364-503X, 1471-2962. doi:10.1098/rsta.2021.0307.

[28] William Waites, Carl A. B. Pearson, Katherine M. Gaskell, Thomas House, Lorenzo Pellis, Marina Johnson, Victoria Gould, Adam Hunt, Neil R. H. Stone, Ben Kasstan, Tracey Chantler, Sham Lal, Chrissy H. Roberts, David Goldblatt, CMMID COVID-19 Working Group, Kaja Abbas, Sam Abbott, Katherine E. Atkins, Rosanna C. Barnard, Nikos I. Bosse, Oliver Brady, Yung-Wai Desmond Chan, Lloyd A. C. Chapman, Samuel Clifford, Nicholas G. Davies, W. John Edmunds, Rosalind M. Eggo, Akira Endo, Emilie Finch, Stefan Flasche, Anna M. Foss, Sebastian Funk, Hamish P. Gibbs, Amy Gimma, Joel Hellewell, David Hodgson, Stéphane Hué, Yalda Jafari, Christopher I. Jarvis, Mark Jit, Thibaut Jombart, Petra Klepac, Gwenan M. Knight, Mihaly Koltai, Fabienne Krauer, Adam J. Kucharski, Yang Liu, Rachel Lowe, Ciara V. McCarthy, Sophie R. Meakin, Graham Medley, Paul Mee, James D. Munday, Emily S. Nightingale, Kathleen O'Reilly, Carl A. B. Pearson, Kiesha Prem, Simon R. Procter, Rachael Pung, Matthew Quaife, Billy J. Quilty, Alicia Rosello, Timothy W. Russell, Frank G. Sandmann, Fiona Yueqian Sun, Damien C. Tully, C. Julian Villabona-Arenas, William Waites, Naomi R. Waterlow, Kerry L. M. Wong, Kevin van Zandvoort, Michael Marks, and Rosalind M. Eggo. Transmission dynamics of SARS-CoV-2 in a strictly-Orthodox Jewish community in the UK. *Scientific Reports*, 12(1):8550, December 2022. ISSN 2045-2322. doi:10.1038/s41598-022-12517-6.

[29] Robert H Chisholm, Patricia T Campbell, Yue Wu, Steven Y C Tong, Jodie McVernon, and Nicholas Geard. Implications of asymptomatic carriers for infectious disease transmission and control. *Royal Society Open Science*, 5(2):172341, 2018. doi:10.1098/rsos.172341.

[30] Sang Woo Park and Benjamin M Bolker. A note on observation processes in epidemic models. *Bulletin of Mathematical Biology*, 82(3):36, 2020. doi:10.1007/s11538-020-00713-2.

[31] Eunha Shim, Zhilan Feng, and Carlos Castillo-Chavez. Differential impact of sickle cell trait on symptomatic and asymptomatic malaria. *Mathematical Biosciences and Engineering*, 9(4):877–898,

2012. doi:10.3934/mbe.2012.9.877.

[32] Pierre Boutillier, Mutaamba Maasha, Xing Li, Héctor F. Medina-Abarca, Jean Krivine, Jérôme Feret, Ioana Cristescu, Angus G. Forbes, and Walter Fontana. The Kappa platform for rule-based modeling. *Bioinformatics*, 34(13):i583–i592, July 2018. ISSN 1367-4803. doi:10.1093/bioinformatics/bty272. Publisher: Oxford Academic.

[33] Pierre Boutillier, Jérôme Feret, Jean Krivine, and Walter Fontana. Kasim: a stochastic simulator for rule-based models written in kappa. https://github.com/Kappa-Dev/KaSim, 2024. Accessed February 2026.

[34] Banu Aygun and Isaac Odame. A global perspective on sickle cell disease. *Pediatric Blood and Cancer*, pages 386–390, 2012.

[35] Lourdes Esteva, Abba B Gumel, and Cruz Vargas De LeóN. Qualitative study of transmission dynamics of drug-resistant malaria. *Mathematical and Computer Modelling*, 50(3-4):611–630, 2009.

[36] Joan L Aron and Robert M May. The population dynamics of malaria. In Roy Anderson, editor, *The population dynamics of infectious diseases: theory and applications*, pages 139–179. Springer, 1982.

[37] Thomas S Churcher, Robert E Sinden, Nick J Edwards, Ian D Poulton, Thomas W Rampling, Patrick M Brock, Jamie T Griffin, Leanna M Upton, Sara E Zakutansky, Katarzyna A Sala, et al. Probability of transmission of malaria from mosquito to human is regulated by mosquito parasite density in naive and vaccinated hosts. *PLOS Pathogens*, 13(1):e1006108, 2017.

[38] Thomas S Churcher, Jean-Francois Trape, and Anna Cohuet. Human-to-mosquito transmission efficiency increases as malaria is controlled. *Nature Communications*, 6(1):6054, 2015.

[39] CDC. Malaria. https://www.cdc.gov/malaria/about/biology/index.html, 2020. Accessed November 16, 2023.

[40] World Health Organization. Malaria control in complex emergencies: an inter-agency field handbook. Technical report, World Health Organization, 2005.

[41] Katya Galactionova, Fabrizio Tediosi, Don de Savigny, Thomas Smith, and Marcel Tanner. Effective coverage and systems effectiveness for malaria case management in sub-saharan african countries. *PLOS ONE*, pages 1–21, 2015.

[42] Aleisha R Brock, Joshua V Ross, Scott Greenhalgh, et al. Modelling the impact of antimalarial quality on the transmission of sulfadoxine-pyrimethamine resistance in plasmodium falciparum. *Infectious Disease Modelling*, pages 161–187, 2017.

[43] Chiara Andolina, Wouter Graumans, Moussa W Guelbeogo, et al. Plasmodium falciparum gametocyte carriage in longitudinally monitored incident infections is associated with duration of infection and human host factors. *Scientific Reports*, 13(7072):7042, 2023. doi:10.1038/s41598-023-33657-3.

[44] Jaline Gerardin, Alfred L Ouédraogo, Kevin A McCarthy, Philip A Eckhoff, and Edward A Wenger. Characterization of the infectious reservoir of malaria with an agent-based model calibrated to age-stratified parasite densities and infectiousness. *Malaria Journal*, 14:231, 2015. doi:10.1186/s12936-015-0751-y.